arXiv:1704.08468v1 [cs.DS] 27 Apr 2017

Linear-Size Hopsets with Small Hopbound, and Distributed Routing with Low Memory

Michael Elkin*[1] and Ofer Neiman†[1]

[1]Department of Computer Science, Ben-Gurion University of the Negev, Beer-Sheva, Israel.
Email: {elkinm,neimano}@cs.bgu.ac.il

Abstract

For a positive parameter β, the *β-bounded distance* between a pair of vertices u, v in a weighted undirected graph $G = (V, E, \omega)$ is the length of the shortest $u - v$ path in G with at most β edges, aka *hops*. For β as above and $\epsilon > 0$, a *(β, ϵ)-hopset* of $G = (V, E, \omega)$ is a graph $G' = (V, H, \omega_H)$ on the same vertex set, such that all distances in G are $(1+\epsilon)$-approximated by β-bounded distances in $G \cup G'$.

Hopsets are a fundamental graph-theoretic and graph-algorithmic construct, and they are widely used for distance-related problems in a variety of computational settings. Currently existing constructions of hopsets produce hopsets either with $\Omega(n \log n)$ edges, or with a hopbound $n^{\Omega(1)}$. In this paper we devise a construction of *linear-size* hopsets with hopbound (ignoring the dependence on ϵ) $(\log n)^{\log^{(3)} n + O(1)}$. This improves the previous bound almost exponentially.

We also devise efficient implementations of our construction in PRAM and distributed settings. The only existing PRAM algorithm [EN16a] for computing hopsets with a constant (i.e., independent of n) hopbound requires $n^{\Omega(1)}$ time. We devise a PRAM algorithm with polylogarithmic running time for computing hopsets with a constant hopbound, i.e., our running time is *exponentially* better than the previous one. Moreover, these hopsets are also significantly sparser than their counterparts from [EN16a].

We use our hopsets to devise a distributed routing scheme that exhibits near-optimal tradeoff between individual memory requirement $\tilde{O}(n^{1/k})$ of vertices throughout preprocessing and routing phases of the algorithm, and stretch $O(k)$, along with a near-optimal construction time $\approx D + n^{1/2+1/k}$, where D is the hop-diameter of the input graph. Previous distributed routing algorithms either suffered from a prohibitively large memory requirement $\Omega(\sqrt{n})$, or had a near-linear construction time, even on graphs with small hop-diameter D.

*This research was supported by the ISF grant No. (724/15).
†Supported in part by ISF grant No. (523/12) and by BSF grant No. 2015813.

1 Introduction

1.1 Hopsets

Consider a weighted undirected graph $G = (V, E, \omega)$. Consider another graph $G_H = (V, H, \omega_H)$ on the same vertex set, that satisfies that for every $(u, v) \in H$, $\omega_H(u, v) \geq d_G(u, v)$, where $d_G(u, v)$ stands for the *distance* between u and v in G. For a positive integer parameter β, and a positive parameter $\epsilon > 0$, the graph G_H is called a (β, ϵ)-*hopset* of G, if for every pair of vertices $u, v \in V$, the β-*bounded distance* $d^{(\beta)}_{G \cup G_H}(u, v)$ between u and v in $G \cup G_H$ is within a factor $1 + \epsilon$ from the distance $d_G(u, v)$ between these vertices in G, i.e., $d_G(u, v) \leq d^{(\beta)}_{G \cup G_H}(u, w) \leq (1 + \epsilon) d_G(u, v)$. The β-*bounded distance* between u and v in G is the length of the shortest β-*bounded* u-v-*path* $\Pi_{u,v}$ in G, i.e., of a shortest path $\Pi_{u,v}$ with at most β hops/edges.

Hopsets constitute an important algorithmic and combinatorial object, and they are extremely useful for approximate distance computations in a large variety of computational settings, including the distributed model [Nan14, HKN16, EN16a, EN16b, Elk17], parallel (PRAM) model [KS97, SS99, Coh00, MPVX15, EN16a], streaming model [Nan14, HKN16, EN16a, Elk17], dynamic setting [Ber09, HKN14], and for routing [LP15, EN16a]. The notion of hopsets was coined in a seminal paper of Cohen [Coh00]. (Though some first implicit constructions appeared a little bit earlier [UY91, KS97, SS99, Coh97].

In [Coh00], Cohen also devised landmark constructions of hopsets. Specifically, she showed that for any parameters $\kappa = 1, 2, \ldots$, and $\epsilon > 0$, and any n-vertex graph G, there exists a (β, ϵ)-hopset with $O(n^{1+1/\kappa} \cdot \log n)$ edges, with $\beta = O\left(\frac{\log n}{\epsilon}\right)^{O(\log \kappa)}$. Moreover, she showed that hopsets with comparable attributes can be efficiently constructed in the centralized and parallel settings. Specifically, in the centralized setting her algorithm takes an additional parameter $\rho > 0$, and constructs a hopset of size $O(n^{1+1/\kappa} \cdot \log n)$ edges in $O(|E| n^{\rho})$ time, with hopbound $\beta = O\left(\frac{\log n}{\epsilon}\right)^{\frac{O(\log \kappa)}{\rho}}$. In the PRAM model, her hopset has size $O(n^{1+1/\kappa}) \cdot O(\log n)^{O(\frac{\log \kappa}{\rho})}$, and β is as above, and it is constructed in roughly $O(\beta)$ (i.e., polylogarithmic time), and with $O(|E| \cdot n^{\rho})$ work.

Cohen [Coh00] also raised the open question of existence and efficient constructability of hopsets with better attributes; she called it an "intriguing research problem". In the two decades that passed since Cohen's work [Coh00], numerous algorithms for constructing hopsets in various settings were devised [Ber09, HKN14, Nan14, MPVX15, EN16a]. The hopsets of [Ber09, HKN14, Nan14, HKN16] are no better than those of [Coh00] in terms of their attributes. (But they are constructed in settings to which Cohen's algorithm is not known to apply.) The algorithm of [MPVX15] builds hopsets of size $O(n)$, but with a large hopbound $\beta = n^{\Omega(1)}$. In [EN16a], the current authors showed that there always exist (β, ϵ)-hopsets of size $O(n^{1+1/\kappa} \cdot \log n)$, with *constant* (i.e., independent of n) hopbound $\beta = O(\frac{\log \kappa}{\epsilon})^{\log \kappa + O(1)}$. Abboud et al. [ABP17] showed a lower bound $\beta = \Omega(\frac{1}{\epsilon \log \kappa})^{\log \kappa}$ on the hopbound of hopsets with size $O(n^{1+1/\kappa})$.

In the PRAM model, [EN16a] showed two results. First, that for parameters $\epsilon, \rho, \zeta > 0$ and $\kappa = 1, 2, \ldots$, hopsets with constant hopbound $\beta = O\left(\frac{\log \kappa + 1/\rho}{\epsilon \zeta}\right)^{\log \kappa + 2/\rho + O(1)}$ can be constructed in time $O(n^{\zeta} \cdot \beta)$, using $O(|E| \cdot n^{\rho + \zeta})$ work. Second, [EN16a] devised a PRAM algorithm with polylogarithmic time $O\left(\frac{\log n}{\epsilon}\right)^{\log \kappa + 1/\rho + O(1)}$, albeit with a polylogarithmic β roughly equal to the running time. The hopset's size is $O(n^{1+1/\kappa} \cdot \log n)$ in the second result as well.

These results of [EN16a] strictly outperformed the longstanding tradeoff of [Coh00] in all regimes, and proved existence of hopsets with constant hopbound. However, they left a significant room for improvement.

First, the hopsets of [EN16a] have $\Omega(n \log n)$ edges in all regimes. As a result, the only currently existing sparser hopsets [UY91, KS97, SS99, MPVX15] have hopbound of $n^{\Omega(1)}$. Hence it was left open in [EN16a] if hopsets with size $o(n \log n)$ and hopbound $n^{o(1)}$ exist. Second, the question whether hopsets with constant hopbound can be constructed in polylogarithmic PRAM time was left open in [EN16a]. Indeed, the hopsets' algorithm of [EN16a] for constructing such hopsets requires $n^{\Omega(1)}$ PRAM time.

In this paper we answer both these questions in the affirmative. Specifically, we show that for any $\kappa = 1, 2, \ldots$, there exists a $(\beta(\kappa, \epsilon), \epsilon)$-hopset, for all $\epsilon > 0$ *simultaneously*, with size $O(n^{1+1/\kappa})$ and $\beta = \beta(\kappa, \epsilon) = O\left(\frac{\log \kappa}{\epsilon}\right)^{\log \kappa + O(1)}$. In particular, by setting $\kappa = \log n$, we obtain a *linear-size* hopset, and its hopbound is $\beta = O\left(\frac{\log \log n}{\epsilon}\right)^{\log \log n + O(1)}$. This is an almost exponential improvement of the previously best known upper bound (due to [MPVX15]) on the hopbound of linear-size hopsets.

Second, in the PRAM setting, for any $\kappa = 1, 2, \ldots$, $\epsilon > 0$, and $\rho > 0$, our algorithm constructs hopsets of size $O(n^{1+1/\kappa} \cdot \log^* n)$, *constant hopbound* $\beta = O\left(\frac{(\log \kappa + 1/\rho)^2}{\epsilon}\right)^{\log \kappa + 1/\rho + O(1)}$, in *polylogarithmic* time $O((\log n)/\epsilon)^{\log \kappa + 1/\rho + O(1)}$, using work $O(|E| \cdot n^{\rho})$. This is an exponential improvement of the parallel running time of the previously best-known algorithm for constructing hopsets with constant hop-bound due to [EN16a]. We can also shave the $\log^* n$ factor in the size, that allows for a linear-size hopset, but then β grows to be roughly the running time. See Table 1 for a concise comparison between existing and new results concerning hopsets in the PRAM model.

Our algorithm also provides improved results for constructing hopsets in distributed CONGEST and Congested Clique models (see Section 3 for definition of these models). In all these models, our algorithm constructs linear-size hopsets. Also, the running time of our algorithms in all these models is purely combinatorial, i.e., it does not depend on the aspect ratio Λ of the graph. [1] In contrast, previous algorithms [Nan14, HKN16, EN16a] for constructing hopsets in the CONGEST model all have running time proportional to $\log \Lambda$.

Reference	Size	β = Hopbound	Time	Work
[KS97, SS99]	$O(n)$	$O(\sqrt{n})$	$O(\sqrt{n} \log n)$	$O(\lvert E\rvert \cdot \sqrt{n})$
[MPVX15]	$O(n)$	$O(n^{\frac{4+\alpha}{4+2\alpha}})$	$O(n^{\frac{4+\alpha}{4+2\alpha}})$	$O(\lvert E\rvert \cdot \log^{3+\alpha} n)$
	$O(n)$	$O(n^{\alpha})$ $(\alpha \geq \Omega(1))$	$O(n^{\alpha})$	$O(\lvert E\rvert \cdot \log^{O(1/\alpha)} n)$
[Coh00]	$n^{1+1/\kappa} \cdot (\log n)^{O(\frac{\log \kappa}{\rho})}$	$(\log n)^{O(\frac{\log \kappa}{\rho})}$	$(\log n)^{O(\frac{\log \kappa}{\rho})}$	$O(\lvert E\rvert \cdot n^{\rho})$
[EN16a]	$O(n^{1+\frac{1}{\kappa}} \cdot \log n)$	$(\log n)^{\log \kappa + \frac{1}{\rho} + O(1)}$	$O(\beta)$	$O(\lvert E\rvert \cdot n^{\rho})$
	$O(n^{1+\frac{1}{\kappa}} \cdot \log n)$	$\left(\frac{\log \kappa + \frac{1}{\rho}}{\zeta}\right)^{\log \kappa + O(\frac{1}{\rho})}$	$O(n^{\zeta} \cdot \beta)$	$O(\lvert E\rvert \cdot n^{\rho+\zeta})$
This paper	$O(n^{1+\frac{1}{\kappa}})$	$(\log n)^{\log \kappa + \frac{1}{\rho} + O(1)}$	$O(\beta)$	$O(\lvert E\rvert \cdot n^{\rho})$
	$O(n^{1+\frac{1}{\kappa}} \cdot \log^* n)$	$\left(k + \frac{1}{\rho}\right)^{O(\log \kappa + \frac{1}{\rho})}$	$(\log n)^{\log \kappa + \frac{1}{\rho} + O(1)}$	$O(\lvert E\rvert \cdot n^{\rho})$

Table 1: Comparison between (β, ϵ)-hopsets in the PRAM model (neglecting the dependency on ϵ). The hopsets of [KS97, SS99] provide exact distances.

[1] The aspect ratio Λ of a graph G is given by $\Lambda = \frac{\max_{u,v \in V} d_G(u,v)}{\min_{u,v \in V, u \neq v} d_G(u,v)}$.

1.2 Distributed Routing with Small Memory

The main application of our novel hopsets' construction is to the problem of distributed construction of compact routing schemes. A routing scheme has two main phases: the preprocessing phase, and the routing phase. In the preprocessing phase, each vertex is assigned a routing table and a routing label.[2] In the routing phase, a vertex u gets a message M with a short header $Header(M)$ and with a destination label $Label(v)$ of a vertex v, and based on its routing table $Table(u)$, on $Label(v)$, and on $Header(M)$, the vertex u decides to which neighbor $x \in \Gamma(u)$ to forward the message M, and which header to attach to the message. The *stretch* of a routing scheme is the worst-case ratio between the length of a path on which a message M travels, and the graph distance between the message's origin and destination.

Due to its both theoretical and practical appeal, routing is a central problem in distributed graph algorithms [PU89, ABNLP90, TZ01b, Cow01, EGP03, GP03, AGM04, Che13]. A landmark routing scheme was devised in [TZ01b]. For an integer $k \geq 1$, the stretch of their scheme is $4k-5$, the tables are of size $O(n^{1/k}\log^{1-1/k} n)$, the labels are of size $O(k \log n)$, and the headers are of size $O(\log n)$. Chechik [Che13] improved this result, and devised a scheme with stretch $3.68k$, and other parameters like in [TZ01b].

An active thread of research [ABNLP90, AP92, LP13, GGHI13, LP15, EN16b] focuses on efficient implementation of the *preprocessing* phase of routing in the distributed CONGEST model, i.e., computing compact tables and short labels that enable for future low-stretch routing. This problem was raised in a seminal paper by Awerbuch, Bar-Noy, Linial and Peleg [ABNLP90], who devised a routing scheme with stretch $2^{O(k)}$, overall memory requirement $\tilde{O}(n^{1+1/k})$,[3] individual memory requirement for a vertex v of $\tilde{O}(\deg(v)+n^{1/k})$, and construction time $\tilde{O}(n^{1+1/k})$ (in the CONGEST model). The "individual memory requirement" parameter encapsulates the routing tables and labels, and the memory used while computing the tables and labels.

Lenzen and Patt-Shamir [LP15] devised a distributed routing scheme (based on [TZ01b]) with stretch $4k-3+o(1)$, tables of size $O(n^{1/k}\log n)$, labels of size $O(k \log n)$, individual memory requirement of $\tilde{O}(n^{1/k})$, and construction time $\tilde{O}(S+n^{1/k})$, where S is the *shortest-path diameter* of the input graph G, i.e., the maximum number of hops in a shortest path between a pair of vertices in G. Though S is often much smaller than n, it is desirable to evaluate complexity measures of distributed algorithms in terms of n and D, where D is the *hop-diameter* of G, defined as the maximum distance between a pair of vertices u, v in the underlying unweighted graph of G. Typically, we have $D \ll S \ll n$, and it is always the case that $D \leq S \leq n$. (See Peleg's book [Pel00] for a comprehensive discussion.)

Lenzen and Patt-Shamir [LP13] also devised a routing scheme with tables of size $\tilde{O}(n^{1/2+1/k})$, labels of size $O(\log n \cdot \log k)$, stretch at most $O(k \log k)$, and has running time of $\tilde{O}(n^{1/2+1/k}+D)$ rounds. They (based on [SHK+12]) also showed a lower bound of $\tilde{\Omega}(D+\sqrt{n})$ on the time required to construct a routing scheme. In a follow-up paper, [LP15] showed how to improve the stretch of the above scheme to $O(k)$. The main drawback of this result is the prohibitively large size of the routing tables. (The individual memory requirement is consequently prohibitively large as well.) They also exhibited a different tradeoff, that overcame the issue of large routing tables. They devised an algorithm that produced routing tables of size $O(n^{1/k}\cdot\log^2 n)$, labels of size $O(k\log^2 n)$ and stretch $4k-3+o(1)$, albeit with sub-optimal running time $\tilde{O}(\min\{(nD)^{1/2}n^{1/k}, n^{2/3}+D\})\cdot\log\Lambda$, and no guarantee on the individual memory requirement during the preprocessing phase. In [EN16b], the current authors improved the bounds of [LP13, LP15]. In

[2] In this paper we only consider *labeled* or *name-dependent* routing, in which vertices are assigned labels by the scheme. There is also a large body of literature on name-independent routing schemes; cf. [AGM+08] and the references therein. However, a lower bound [LP13] shows that constructing a name-independent routing scheme with stretch ρ requires $\tilde{\Omega}(n/\rho^2)$ time in the CONGEST model. See also [GGHI13] for lower bounds on the *communication* complexity of the preprocessing phase of distributed routing.

[3] $\tilde{O}(f(n))$-notation hides polylogarithmic in $f(n)$ factors.

the current state-of-the-art scheme [EN16b], the stretch is $4k-5+o(1)$, the tables and labels are of the same size as in [LP13, LP15] (i.e., $O(n^{1/k}\log^2 n)$ and $O(k\log^2 n)$, respectively), the construction time is $O((n^{1/2+1/k}+D)\cdot\min\{(\log n)^{O(k)},2^{\tilde{O}(\sqrt{\log n})}\}\cdot\log\Lambda$. (A similar, though slightly weaker, result was achieved by [LPP16].) Still there is no meaningful guarantee on the individual memory requirement in the preprocessing phase. See Table 2 for a concise summary of existing bounds, and a comparison with our new results.

To summarize, all currently existing distributed routing algorithms with nearly-optimal running time $\approx D+n^{1/2+1/k}$ suffer from three issues. First, they provide no meaningful guarantee on individual memory requirement on vertices in the preprocessing phase; second, their preprocessing time is not purely combinatorial, but rather depends linearly on $\log\Lambda$; and third, their tables and labels sizes are roughly $O(\log n)$ off from the respective tables and labels' sizes of Thorup-Zwick's sequential construction [TZ01b].

The issue of individual memory requirement was indeed explicitly raised by Lenzen [Len16] in a private communication with the authors. He wrote (the stress on "during" is in the origin):

"*One annoying thing about this is that there is a huge amount of storage required *during* the construction. It seems odd that the nodes can hold only small tables, but should have large memory during the construction. I think it's an interesting question whether we can have a good construction using $\tilde{O}(n^{1/k})$ memory only. A hop set may be the wrong option for this, because reflecting the distance structure of the skeleton accurately cannot be done by a sparse graph; on the other hand, maybe there's some distributed representation cleverly distributing the information over the graph nodes?*"

Based on our novel hopsets' construction, we devise an algorithm that addresses all these issues. Specifically, the stretch of our scheme is $4k-5+o(1)$, the sizes of tables and labels essentially match the respective sizes of Thorup-Zwick's construction, i.e., they are $O(n^{1/k}\log n)$ and $O(k\log n)$, respectively. Our construction time is $(n^{1/2+1/k}+D)\cdot(\log n)^{O(k)}$, i.e., it is purely combinatorial. Most importantly, the individual memory requirement is at most $\tilde{O}(n^{1/k})$. Moreover, we can reduce the running time to $(n^{1/2+1/k}+D)\cdot\min\{(\log n)^{O(k)},2^{\tilde{O}(\sqrt{\log n})}\}$, while the individual memory increases slightly to $\max\{\tilde{O}(n^{1/k}),2^{\tilde{O}(\sqrt{\log n})}\}$. In particular, we can have a polylogarithmic individual memory requirement and construction time $O((n^{1/2}+D)\cdot n^{\epsilon})$, for an arbitrarily small constant $\epsilon>0$.

Distributed Tree Routing: An important ingredient in the existing distributed routing schemes [LP15, EN16b] for general graphs, and in our new routing scheme, is a distributed tree routing scheme. Thorup and Zwick [TZ01b] showed that with routing tables of size $O(1)$ and labels of size $O(\log n)$, one can have an *exact* (i.e., no stretch) tree routing. [LP15, EN16b] showed that in $\tilde{O}(D+\sqrt{n})$ time, one can construct exact tree routing with tables and labels of size $O(\log n)$ and $O(\log^2 n)$, respectively, i.e., there is an overhead of $\log n$ in both parameters with respect to Thorup-Zwick's sequential construction. In this paper we improve this result, and devise a $\tilde{O}(D+\sqrt{n})$-time algorithm that constructs tree-routing tables and labels of sizes that match the sequential construction of Thorup and Zwick, i.e., of size $O(1)$ and $O(\log n)$, respectively. Moreover, if one is interested in a scheme that always routes via the root of the tree, as is the case in the application to routing in general networks, then our algorithm for constructing tables and labels that supports this requires only a small ($O(\log n)$) individual memory in each vertex.

1.3 Technical Overview

Cohen's algorithm [Coh00] is based on a subroutine for constructing pairwise covers [Coh93, ABCP93], i.e., collections of small-radii clusters with small maximum overlap (no vertex belongs to too many clusters). The

Reference	Number of Rounds	Table size	Label size	Stretch	Memory per vertex
[ABNLP90]	$O(n^{1+\frac{1}{k}})$	$O(n^{\frac{1}{k}} \cdot \log n)$	$O(k \log n)$	$2 \cdot 3^k - 1$	$\tilde{O}(\deg(v)+n^{\frac{1}{k}})$
[TZ01b, Che13]	$O(n^{1+\frac{1}{k}})$	$O(n^{\frac{1}{k}} \cdot \log n)$	$O(k \log n)$	$3.68k$	$O(n^{\frac{1}{k}} \cdot \log n)$
[LP13, LP15]	$\tilde{O}(n^{\frac{1}{2}+\frac{1}{4k}}+D)$	$\tilde{O}(n^{\frac{1}{2}+\frac{1}{4k}})$	$O(\log n)$	$6k-1+o(1)$	$\tilde{O}(n^{\frac{1}{2}+\frac{1}{4k}})$
[LP15]	$\tilde{O}(S+n^{\frac{1}{k}})$	$O(n^{\frac{1}{k}} \cdot \log n)$	$O(k \log n)$	$4k-3$	$O(n^{\frac{1}{k}} \cdot \log n)$
	$\tilde{O}(\min\{(nD)^{\frac{1}{2}} \cdot n^{\frac{1}{k}}, n^{\frac{2}{3}+\frac{2}{3k}}+D\})$	$O(n^{\frac{1}{k}} \cdot \log^2 n)$	$O(k \log^2 n)$	$4k-3+o(1)$	$\tilde{O}(n^{\frac{1}{2}})$
[LPP16]	$(n^{\frac{1}{2}+\frac{1}{k}}+D) \cdot 2^{O(\sqrt{\log n})}$	$O(n^{\frac{1}{k}} \cdot \log^2 n)$	$O(k \log^2 n)$	$4k-3+o(1)$	$\tilde{O}(n^{\frac{1}{2}})$
[EN16b]	$(n^{\frac{1}{2}+\frac{1}{k}}+D) \cdot \beta$	$O(n^{\frac{1}{k}} \cdot \log^2 n)$	$O(k \log^2 n)$	$4k-5+o(1)$	$\tilde{O}(n^{\frac{1}{2}})$
This paper	$(n^{\frac{1}{2}+\frac{1}{k}}+D) \cdot (\log n)^{O(k)}$	$O(n^{\frac{1}{k}} \cdot \log n)$	$O(k \log n)$	$4k-5+o(1)$	$\tilde{O}(n^{\frac{1}{k}})$
	$(n^{\frac{1}{2}+\frac{1}{k}}+D) \cdot 2^{O(\sqrt{\log n})}$	$O(n^{\frac{1}{k}} \cdot \log n)$	$O(k \log n)$	$4k-5+o(1)$	$2^{O(\sqrt{\log n})}$

Table 2: Comparison of compact routing schemes for graphs with n vertices, m edges, hop-diameter D, and shortest path diameter S. Denote $\beta = \min\{(\log n)^{O(k)}, 2^{\tilde{O}(\sqrt{\log n})}\}$.

algorithm is a top-down recursive procedure: it interconnects large clusters of the cover via hopset edges, and recurses in small clusters. To keep the overall overlap of all recursion levels in check, Cohen used the radius parameter $O(\log n)$ for the covers. This resulted in a hop-bound, which is at least polylogarithmic in n. Cohen's hopset is also built separately for each distance scale $[2^i, 2^{i+1})$, $i = 0, 1, \ldots, \log \Lambda$, and the ultimate hopset is the union of all these single-scale hopsets.

The hopset's construction of [EN16a], (due to the current authors), also builds a single-scale hopset for each distance scale, and then takes their union as an ultimate hopset. The construction of single-scale hopsets in [EN16a] is based upon ideas from the construction of $(1+\epsilon, \beta)$-spanners of [EP04] for unweighted graphs. It starts with a partition of the vertex set V into singleton clusters $\mathcal{P}_0 = \{\{v\} \mid v \in V\}$, and alternates superclustering and interconnection steps. In a superclustering step some of the clusters of the current partition are merged into larger clusters (of $\mathcal{P}_{i+1}$), while the other clusters are interconnected with one another via hopset edges.

This approach (of [EN16a]) enabled us to prove existence of hopsets with constant hop-bound, but it appears to be uncapable of producing hopsets of size $o(n \log n)$. Indeed, even if each single-scale hopset is of linear size (which is indeed the case in [EN16a]), their union is doomed to be of size $\Omega(n \log n)$. Moreover, in parallel and distributed settings, one produces a hopset $i+1$ based upon a hopset of scale i. This results in *accumulation of stretch* from $1+\epsilon$ to $(1+\epsilon)^{\log n}$. To alleviate this issue, one needs to rescale ϵ. However, then the hopbound grows from constant to polylogarithmic. To get around this, [EN16a] used a smaller number of scales, and this indeed enables [EN16a] to construct hopsets with constant hopbound in these settings, albeit the running time becomes proportional to the ratio between consequent scales, i.e., it becomes $n^{\Omega(1)}$.

The closest to our current construction of hopsets is the line of research of [Ber09, HKN14, HKN16, Nan14], which is based on a construction of distance oracles due to Thorup and Zwick [TZ01a]. To construct their oracles, [TZ01a] used a hierarchy of sets $V = A_0 \supseteq A_1 \supseteq \ldots A_{\kappa-1} \supseteq A_\kappa = \emptyset$, where each vertex of A_i, for all $i = 0, 1, \ldots, \kappa-2$, is sampled with probability $n^{-1/\kappa}$ for the inclusion into A_{i+1}. For each vertex $u \in V$, [TZ01a] defined for every $i = 0, 1, \ldots, \kappa-1$, the ith *pivot* $p_i(u)$ to be its closest vertex in A_i, and the ith *bunch* $B_i(u) = \{v \mid d_G(u,v) < d_G(u, A_{i+1})\} \cup \{p_{i+1}(u)\}$, (for $i = \kappa$, let $\{p_\kappa(u)\} = \emptyset$), and the entire bunch $B(u) = \bigcup_{i=0}^{\kappa-1} B_i(u)$. They also defined the dual sets, *clusters*, $C(v) = \{u \mid v \in B(u)\}$. Bernstein and others [Ber09, HKN14, Nan14, HKN16] used this construction with $\kappa = \tilde{\Theta}(\sqrt{\log n})$, and built Thorup-Zwick clusters with respect to $2^{\tilde{O}(\sqrt{\log n})}$-bounded distances. As a result, they obtained a so-called $2^{\tilde{O}(\sqrt{\log n})}$-*bounded hopset*, i.e., a hopset which takes care only of pairs $u, v \in V$ of vertices that

admit a $2^{\tilde{O}(\sqrt{\log n})}$-bounded shortest path. They then used this bounded hopset in a certain recursive fashion (see the so-called *hop reduction* of Nanongkai [Nan14]), to obtain their ultimate hopset.

Thorup-Zwick's construction with $\kappa = \tilde{\Theta}(\sqrt{\log n})$ alone introduces into the hopset $n \cdot 2^{\tilde{\Theta}(\sqrt{\log n})}$ edges, and thus, such a hopset cannot be very sparse. In addition, the recursive application of the hop reduction technique results in a hopbound of $2^{\tilde{\Omega}(\sqrt{\log n})}$.

Our construction of hopsets is based upon a construction of Thorup-Zwick's emulators[4], from a different paper by Thorup and Zwick [TZ06]. Specifically, to obtain the hierarchy $V = A_0 \supseteq A_1 \supseteq \ldots A_{\log\kappa - 1} \supseteq A_{\log\kappa} = \emptyset$, one samples each vertex of A_i, for $i = 0, 1, \ldots, \log\kappa - 2$, with probability roughly $n^{-2^i/\kappa}$ for inclusion in A_{i+1}. Then one defines the bunch of a vertex $u \in A_i$ as $B(u) = \{v \in A_i \mid d_G(v,u) < d_G(v, A_{i+1})\} \cup \{p_{i+1}(u)\}$, and sets

$$H = \bigcup_{u \in V} \{(u,v) \mid v \in B(u)\} . \tag{1}$$

For *unweighted* graphs G, [TZ06] showed that H given by (1) is an additive emulator with stretch $\alpha(d) = O(\log\kappa \cdot d^{1-1/(\log\kappa - 1)})$ and $O(\log\kappa \cdot n^{1+1/\kappa})$ edges. By a different proof argument, we show that the *very same* construction provides also a (β, ϵ)-hopset of the same size and with $\beta = O\left(\frac{\log\kappa}{\epsilon}\right)^{\log\kappa - 1}$, for all $\epsilon > 0$ *simultaneously*. Moreover, by adjusting the sampling probabilities, we also shave the $\log\kappa$ factor from both the hopset's and the emulator's size, while increasing the exponent of β by 1. (This also gives rise to the first linear-size emulator with sub-linear additive stretch for unweighted graphs.)

As a result, we obtain a construction of hopsets, which is by far simpler than the previous Thorup-Zwick-based constructions of hopsets [Ber09, HKN14, Nan14, HKN16]. As was discussed above, it also provides hopsets with much better parameters, and it is more adaptable to efficient implementation in various computational settings. Our construction is also much simpler than the constructions of [Coh00, EN16a], which are not based on Thorup-Zwick's hierarchy.

Parallel and distributed implementations of our hopset's construction proceed in scales $\ell = 0, 1, \ldots, \log n$, where on scale ℓ the algorithm constructs a 2^ℓ-bounded hopset. This is different from the situation in [Coh00, EN16a, Nan14, HKN16], where scale ℓ takes care of distances in the range $[2^\ell, 2^{\ell+1})$. An important advantage of this is that we no longer need to take the union of all single-scale hopsets into our hopset; rather we just take the largest-scale hopset as our ultimate hopset. This saves a factor of $\log n$ in the size, and enables us to construct *linear-size* hopsets in parallel and distributed settings. The fact that we do not work with distance scales, but rather with hop-distance-scales, makes it possible to avoid the dependence on $\log\Lambda$ in the distributed construction time, and to achieve a purely combinatorial running time. All previous distributed algorithms for constructing approximate hopsets [Nan14, HKN16, EN16a] have running time proportional to $\log\Lambda$. (A distributed construction of an *exact* hopset [Elk17] by the first-named author, however, avoids this dependence too. Alas, it has a much higher running time.)

The fact that the construction's scales are with respect to hop-distances, as opposed to actual distances, enables us to essentially avoid accumulation of error. This is done by the following recursive procedure. First, we build 2^ℓ-bounded hopsets $H^{(\ell)}$, $\ell = 0, 1, \ldots, \log n$, and set $H(1) = H^{(\log n)}$ to be the highest-scale hopset. This process involves accumulation of stretch, and after rescaling the stretch parameter ϵ, the hopset $H(1)$ ends up having polylogarithmic hopbound β_1. Its construction time is roughly β_1, i.e., polylogarithmic as well. Now we add the hopset into the original graph, and recurse on $G \cup H(1)$. Note that

[4] A graph $G' = (V, E', \omega')$ is called a *sublinear-error emulator* of an unweighted graph $G = (V, E)$, if for every pair of vertices $u, v \in V$, we have $d_G(u,v) \leq d_{G'}(u,v) \leq d_G(u,v) + \alpha(d_G(u,v))$ for some sub-linear stretch function α. If G' is a subgraph of G, it is called a sublinear-error *spanner* of G.

now we only need to process $\log \beta_1 \approx \log\log n$ scales, rather than $\log n$ ones. Hence the accumulation of stretch in the resulting hopset $H(2)$ is much more mild than in $H(1)$. As a result, after rescaling the stretch parameter ϵ, the hopbound β_2 of $H(2)$ is roughly $poly(\log\log n)$. By repeating this recursive process for $\log^* n$ iterations, we eventually achieve a hopset with constant hop-bound in parallel polylogarithmic time. As was mentioned above, this dramatically improves the $n^{\Omega(1)}$ parallel time required in [EN16a] to construct a hopset with constant hopbound.

In the context of distributed routing, like in [LP15, EN16b], our hopset H is constructed on top of a *virtual* (aka skeleton) graph $G' = (V', E')$, where V' is a collection of $\approx \sqrt{n}$ vertices, sampled from the original vertex set V independently with probability $\approx n^{-1/2}$. There is an edge $(u', v') \in E'$ iff there is a $\tilde{O}(\sqrt{n})$-bounded $u' - v'$ path in G. However, since we aim to design an algorithm in which vertices employ only a small memory during the preprocessing phase, we cannot afford computing the virtual graph G'. Rather, somewhat surprisingly, we show that the hopset H for G' can be constructed without ever constructing G' itself! We only compute those edges of G', which are required for constructing the hopset H.

Note, however, that unlike a spanner or a low-stretch spanning tree, hopset is always used *in conjunction* with the graph for which it was constructed. In other words, to compute the Thorup-Zwick routing scheme for the virtual graph G', we conduct Bellman-Ford explorations in $G' \cup H$. So, at the first glance, it seems necessary to eventually compute G', for being able to conduct these Bellman-Ford explorations.

We cut this gordian knot by computing only those edges of G' that are really needed for computing either the hopset H or the TZ routing scheme for $G' \cup H$. This turns out to be (typically) a small fraction of edges of G', and those edges can be computed much more efficiently than the entire G', and using much smaller memory. This idea also enables us to compute lengths of these edges of G' *precisely*, as opposed to approximately, as it was done in [LP15, HKN16, EN16b]. This simplifies the analysis of the resulting scheme. We note that the idea of computing a hopset H without first computing the underlying virtual graph G', and conducting Bellman-Ford explorations in $G' \cup H$ without ever computing G' in its entirety appeared in a recent work [Elk17], by the first-named author. [Elk17] constructed an exact hopset of [SS99, Nan14] with a polynomial hopbound. However, the exact hopset is a much simpler structure than the small-hopbound approximate hopset that we construct here. Showing that this idea is applicable for our new approximate small-hopbound hopset is technically substantially more challenging.

Another crucial idea that we employ to guarantee a small individual memory requirement is ensuring that our hopset H has small *arboricity*, i.e., that its edges can be oriented in such a way that every vertex has only a small out-degree. This out-degree is proportional to the ultimate individual memory requirement of our algorithm.

1.4 Organization

Our linear-size hopsets appear in Section 2, and the distributed construction in Section 3.1 for Congested Clique and Section 3.2 for the CONGEST model. The hopsets in the PRAM model are presented in Section 4. Finally, in Section 5 we describe our distributed routing scheme with small memory, and the distributed tree routing in Section 5.1.

2 Linear Size Hopsets

Let $G = (V, E)$ be a weighted graph, and fix a parameter $k \geq 1$. Let $\nu = 1/(2^k - 1)$ (one should think of $\kappa = 1/\nu$). Let $V = A_0 \supseteq A_1 \supseteq \cdots \supseteq A_k = \emptyset$ be a sequence of sets, such that for all $0 \leq i < k - 1$, A_{i+1}

is created by sampling every element from A_i independently with probability $p_i = n^{-2^i \cdot \nu}$.[5] It follows that for $0 \leq i \leq k-1$ we have

$$N_i := \mathbb{E}[|A_i|] = n \cdot \prod_{j=0}^{i-1} p_j = n^{1-(2^i-1)\nu} ,$$

and in particular $N_{k-1} = n^{(1+\nu)/2}$.

For every $0 \leq i \leq k-1$ and every vertex $u \in A_i \setminus A_{i+1}$, define the pivot $p(u) \in A_{i+1}$ as a vertex satisfying $d_G(u, A_{i+1}) = d_G(u, p(u))$ (note $p(u)$ does not exist for $u \in A_{k-1}$), and define the bunch

$$B(u) = \{v \in A_i \ : \ d_G(u,v) < d_G(u, A_{i+1})\} \cup \{p(u)\} .$$

The hopset is created by taking $H = \{(u,v) \ : \ u \in V, v \in B(u)\}$, where the length of the edge (u,v) is set as $d_G(u,v)$. As argued in [TZ01a], for any $0 \leq i \leq k-2$ and $u \in A_i \setminus A_{i+1}$, the size of $B(u)$ is bounded by a random variable sampled from a geometric distribution with parameter p_i (this corresponds to the first vertex of A_i, when ordered by distance to u, that is included in A_{i+1}). Hence $\mathbb{E}[|B(u)|] \leq 1/p_i = n^{2^i \cdot \nu}$. For $u \in A_{k-1}$ we have $\mathbb{E}[|B(u)|] = N_{k-1} = n^{(1+\nu)/2}$. The expected size of the hopset H is at most

$$\sum_{i=0}^{k-2} (N_i \cdot n^{2^i \cdot \nu}) + N_{k-1} \cdot N_{k-1} = k \cdot n^{1+\nu} .$$

The following lemma bounds the number of hops and stretch of H. Recall that $d_G^{(t)}(u,v)$ is the length of the shortest-path between u, v in G that consists of at most t edges.

Lemma 1. *Fix any $0 < \delta < 1/(8k)$ and any $x, y \in V$. Then for every $0 \leq i \leq k-1$, at least one of the following holds:*

1. $d_{G \cup H}^{((3/\delta)^i)}(x,y) \leq (1 + 8\delta i) \cdot d_G(x,y)$.
2. *There exists $z \in A_{i+1}$ such that* $d_{G \cup H}^{((3/\delta)^i)}(x,z) \leq 2 d_G(x,y)$.

Proof. The proof is by induction on i. We start with the basis $i = 0$. If it is the case that $y \in B(x)$, then we added the edge (x,y) to the hopset, i.e. $d_H^{(1)}(x,y) = d_G(x,y)$, and so the first item holds. Otherwise, consider the case that $x \in A_1$: then we can take $z = x$, so the second item holds trivially. The remaining case is that $x \in A_0 \setminus A_1$ and $y \notin B(x)$, so by definition of $B(x)$ we get that $d_G(x,y) \geq d_G(x, A_1)$. By taking $z = p(x)$, there is a single edge between x, z in H of length $d_G(x,z) = d_G(x, A_1) \leq d_G(x,y)$, which satisfies the second item.

Assume the claim holds for i, and we prove for $i+1$. Consider the path $\pi(x,y)$ between x, y in G, and partition it into $J \leq 1/\delta$ segments $\{L_j = [u_j, v_j]\}_{j \in [J]}$, each of length at most $\delta \cdot d_G(x,y)$, and at most $1/\delta$ edges $\{(v_j, u_{j+1})\}_{j \in [J]}$ of G between these segments. This can be done as follows: define $u_1 = x$, and for $j \geq 1$, walk from u_j on $\pi(x,y)$ (towards y) until the point v_j, which is the vertex so that the next edge will take us to distance greater than $\delta \cdot d_G(x,y)$ from u_j (or until we reached y). By definition, $d_G(u_j, v_j) \leq \delta \cdot d_G(x,y)$. Define u_{j+1} to be the neighbor of v_j on $\pi(x,y)$ that is closer to y (if exists). If u_{j+1} does not exist (which can happen only when $v_j = y$) then define $u_{j+1} = y$ and $J = j$. Observe that for all $1 \leq j \leq J-1$, $d_G(u_j, u_{j+1}) > \delta \cdot d_G(x,y)$, so indeed $J \leq 1/\delta$.

[5] Our definition is slightly different than that of [TZ06], which used $p_i = |A_i|/n^{1+\nu}$, but it gives rise to the same expected size of A_i. We use our version since it allows efficient implementation in various models of computation.

We use the induction hypothesis for all pairs (u_j, v_j) with parameter i. Consider first the case that the first item holds for all of them, that is, $d_{G\cup H}^{((3/\delta)^i)}(u_j, v_j) \leq (1 + 8\delta i) \cdot d_G(u_j, v_j)$. Then we take the path in $G \cup H$ that consists of the $(3/\delta)^i$-hops between each pair u_j, v_j, and the edges (v_j, u_{j+1}) of G. Since $(3/\delta)^{i+1} \geq (1/\delta) \cdot (3/\delta)^i + 1/\delta$, we have

$$d_{G\cup H}^{((3/\delta)^{i+1})}(x, y) \leq \sum_{j\in[J]} (d_{G\cup H}^{((3/\delta)^i)}(u_j, v_j) + d_G^{(1)}(v_j, u_{j+1})) \leq (1 + 8\delta i) \cdot d_G(x, y) \; .$$

The second case is that there are pairs (u_j, v_j) for which only the second item holds. Let $l \in [J]$ (resp., $r \in [J]$) be the first (resp., last) index for which the first item does not hold for the pair (u_l, v_l) (resp., (u_r, v_r)). Then there are $z_l, z_r \in A_{i+1}$ such that

$$d_{G\cup H}^{((3/\delta)^i)}(u_l, z_l) \leq 2d_G(u_l, v_l) \text{ and } d_{G\cup H}^{((3/\delta)^i)}(v_r, z_r) \leq 2d_G(u_r, v_r). \tag{2}$$

(Note that we used v_r and not u_r in the second inequality. This can be done since the lemma's assertion holds for the pair (v_r, u_r) as well, and as the first item is symmetric with respect to u_r, v_r, it does not hold for the pair (v_r, u_r) as well.) Consider now the case that $z_r \in B(z_l)$. In this case we added the edge (z_l, z_r) to the hopset, and by the triangle inequality,

$$d_H^{(1)}(z_l, z_r) = d_G(z_l, z_r) \leq d_{G\cup H}^{((3/\delta)^i)}(u_l, z_l) + d_G(u_l, v_r) + d_{G\cup H}^{((3/\delta)^i)}(z_r, v_r) \; . \tag{3}$$

Next, apply the inductive hypothesis on segments $\{L_j\}$ for $j < l$ and $j > r$, and in between use the detour via u_l, z_l, z_r, v_r. Since $l \leq r$, there is at least one segment we skipped, so the total number of hops is bounded by $(1/\delta - 1) \cdot (3/\delta)^i + 1/\delta + 2(3/\delta)^i + 1$. (The additive term of $1/\delta$ accounts for the edges (v_j, u_{j+1}), $0 \leq j \leq J - 1$.) This expression is at most $(3/\delta)^{i+1}$ whenever $\delta < 1/2$. It follows that

$$\begin{aligned}
& d_{G\cup H}^{((3/\delta)^{i+1})}(x, y) \qquad (4)\\
\leq\; & \sum_{j=1}^{l-1} \left[d_{G\cup H}^{((3/\delta)^i)}(u_j, v_j) + d_G^{(1)}(v_j, u_{j+1}) \right] + d_{G\cup H}^{((3/\delta)^i)}(u_l, z_l) + \; d_H^{(1)}(z_l, z_r) \\
& + d_{G\cup H}^{((3/\delta)^i)}(z_r, v_r) + d_G^{(1)}(v_r, u_{r+1}) + \sum_{j=r+1}^{J} \left[d_{G\cup H}^{((3/\delta)^i)}(u_j, v_j) + d_G^{(1)}(v_j, u_{j+1}) \right] \\
\overset{(3)}{\leq}\; & (1 + 8\delta i) d_G(x, u_l) + d_G(u_l, v_r) + (1 + 8\delta i) d_G(v_r, y) \\
& + 2d_{G\cup H}^{((3/\delta)^i)}(u_l, z_l) + 2d_{G\cup H}^{((3/\delta)^i)}(z_r, v_r) \\
\overset{(2)}{\leq}\; & 8\delta \cdot d_G(x, y) + (1 + 8\delta i) d_G(x, y) \\
=\; & (1 + 8\delta(i + 1)) \cdot d_G(x, y) \; .
\end{aligned}$$

This demonstrates item 1 holds in this case. The final case to consider is that $z_r \notin B(z_l)$. Assume first that $z_l \notin A_{i+2}$. Then taking $z = p(z_l) \in A_{i+2}$, the definition of $B(z_l)$ implies that $d_G(z_l, z) \leq d_G(z_l, z_r)$. We

now claim that item 2 holds for such a choice of z. Indeed, since $(z_l, z) \in H$, we have

$$\begin{aligned} d_{G\cup H}^{((3/\delta)^{i+1})}(x,z) &\leq \sum_{j=1}^{l-1}\left[d_{G\cup H}^{((3/\delta)^i)}(u_j,v_j) + d_G^{(1)}(v_j,u_{j+1})\right] + d_{G\cup H}^{((3/\delta)^i)}(u_l,z_l) + d_H^{(1)}(z_l,z) \\ &\leq (1+8\delta i)d_G(x,u_l) + d_{G\cup H}^{((3/\delta)^i)}(u_l,z_l) + d_G(z_l,z_r) \\ &\overset{(3)}{\leq} (1+8\delta i)d_G(x,u_l) + 2d_{G\cup H}^{((3/\delta)^i)}(u_l,z_l) + d_G(u_l,v_r) + d_{G\cup H}^{((3/\delta)^i)}(z_r,v_r) \\ &\overset{(2)}{\leq} (1+8\delta i)d_G(x,v_r) + 6\delta\cdot d_G(x,y) \\ &\leq 2d_G(x,y), \end{aligned}$$

where the last inequality used that $\delta < 1/(8k)$. The case that $z_l \in A_{i+2}$ is simpler, since we may take $z = z_l$. □

Fix any $0 < \epsilon < 1$, and apply the lemma on any pair x, y with $\delta = \epsilon/(8k)$ and $i = k-1$. It must be that the first item holds (since $A_k = \emptyset$). Hence we have that

$$d_{G\cup H}^{(\beta)}(x,y) \leq (1+\epsilon)d_G(x,y)\ ,$$

where the number of hops is given by $\beta = (24k/\epsilon)^{k-1}$. We derive the following theorem.

Theorem 1. *For any weighted graph $G = (V,E)$ on n vertices, and any $k \geq 1$, there exists H of size at most $O(k\cdot n^{1+1/(2^k-1)})$, which is a (β,ϵ)-hopset for any $0 < \epsilon < 1$ with $\beta = O(k/\epsilon)^{k-1}$.*

2.1 Improved Hopset Size

Here we show how to remove the k factor from the hopset size, at the cost of increasing the exponent of β by an additive 1. Note that we may assume w.l.o.g that $k \leq \log\log n - 1$, as for larger values of k, both β and the size of the hopset (which becomes $O(kn)$), grow with k. We will increase the number of sets by 1, and sample $V = A'_0 \supseteq A'_1 \supseteq \cdots \supseteq A'_{k+1} = \emptyset$ using the following probabilities: $p'_i = n^{-2^i\cdot\nu}\cdot 2^{2^i-1}$ (the restriction on k ensures $p'_i < 1$). Now for $0 \leq i \leq k$,

$$N'_i := \mathbb{E}[|A'_i|] = n\cdot\prod_{j=0}^{i-1} p'_j = n^{1-(2^i-1)\nu}\cdot 2^{2^i-i-1}\ ,$$

and in particular $N'_k \leq 2^{2^k-k} \leq n^{1/2}$. The expected size of H becomes at most

$$\sum_{i=0}^{k-1}(N'_i/p'_i) + N'_k\cdot N'_k \leq \sum_{i=0}^{k-1}(n^{1+\nu}/2^i) + n \leq 3n^{1+\nu}\ .$$

The hopset construction and the stretch analysis in Lemma 1 remains essentially the same. There is an additional sampled set now, and thus the exponent of β grows by an additive 1.

Theorem 2. *For any weighted graph $G = (V,E)$ on n vertices and any $k \geq 1$, there exists H of size at most $O(n^{1+1/(2^k-1)})$, which is a (β,ϵ)-hopset for any $0 < \epsilon < 1$ with $\beta = O(k/\epsilon)^k$.*

Since our construction is based on the [TZ06] emulator construction, following their analysis we obtain an emulator with additive stretch that can have *linear* size.

Corollary 2. *For any un-weighted graph $G = (V,E)$ on n vertices and any $k \geq 1$, there exists an emulator H of size at most $O(n^{1+1/(2^k-1)})$, with additive stretch $O(k\cdot d^{1-1/k})$ for pairs at distance d.*

2.2 Efficient Implementation

We consider the construction and notation of Section 2.1, with the slightly stronger assumption $k \leq \log\log n - 2$.

It was (implicitly) shown in [TZ01a] that for any $0 \leq i < k$, the sets $B(u)$ (and the corresponding distances) for all $u \in A'_i \setminus A'_{i+1}$ can be computed in expected time $O(|E| + n\log n)/p'_i$. (In fact, in [TZ01a], the set $B(u)$ contains more vertices, not only those in A'_i. However, we can remove the extra vertices easily.) The running time becomes larger as i grows, and in order to keep it under control, we use the method of [EN16a]: introduce a parameter $2\nu < \rho < 1$, and redefine the probabilities as follows. Set $i_0 = \lfloor \log(\rho/\nu) \rfloor$ and $i_1 = i_0 + 1 + \left\lceil \frac{1}{\rho} \right\rceil$. For $0 \leq i \leq i_0$, let $p'_i = n^{-2^i \cdot \nu} \cdot 2^{2^i - 1}$ as in Section 2.1. Set also $p'_{i_0+1} = n^{-\rho/2}$, and for the remaining levels $i_0 + 2 \leq i \leq i_1$, set $p'_i = n^{-\rho}$. Finally, let $A'_{i_1+1} = \emptyset$. Note that for $0 \leq i \leq i_0 + 1$, we have that

$$N'_i = n \prod_{j=0}^{i-1} p'_j = n^{1-(2^i-1)\nu} \cdot 2^{2^i - i - 1} \,. \tag{5}$$

In particular, using that $\rho/\nu \leq 2^{i_0+1} \leq 2\rho/\nu$, we get

$$N'_{i_0+1} \leq n^{1-(\rho/\nu-1)\nu} \cdot 2^{2\rho/\nu} \leq n^{1+\nu-\rho/2} \,. \tag{6}$$

(The last inequality uses that $k \leq \log\log n - 2$. Thus $1/\nu = (2^k - 1) \leq \log n/4$, and so that $2^{2\rho/\nu} \leq n^{\rho/2}$.) Thus for any $i \geq i_0 + 2$ we see that

$$N'_i = N'_{i_0+1} \prod_{j=i_0+1}^{i-1} p'_j \overset{(6)}{\leq} n^{1+\nu-\rho/2} \cdot n^{-\rho/2} \cdot n^{-(i-1-(i_0+1))\cdot\rho} = n^{1+\nu} \cdot n^{-(i-(i_0+1))\cdot\rho} \,. \tag{7}$$

The expected number of edges inserted until phase $i_0 + 1$ is at most

$$\sum_{i=0}^{i_0} N'_i/p'_i \overset{(5)}{\leq} \sum_{i=0}^{i_0} n^{1-(2^i-1)\nu} \cdot 2^{2^i-i-1} \cdot n^{2^i\cdot\nu}/2^{2^i-1} = \sum_{i=0}^{i_0} n^{1+\nu}/2^i \leq 2n^{1+\nu} \,.$$

The expected number of edges at phase $i_0 + 1$ is bounded by

$$N'_{i_0+1}/p'_{i_0+1} \overset{(6)}{\leq} n^{1+\nu-\rho/2} \cdot n^{\rho/2} = n^{1+\nu} \,.$$

The remaining phases until i_1 introduce at most

$$\sum_{i=i_0+2}^{i_1-1} N'_i/p'_i \overset{(7)}{\leq} \sum_{i=i_0+2}^{i_1-1} n^{1+\nu} \cdot n^{-(i-(i_0+1))\cdot\rho} \cdot n^{\rho} \leq n^{1+\nu} \cdot \sum_{i=0}^{\infty} n^{-i\rho} \leq 2n^{1+\nu} \,,$$

as this summation converges. Finally, since

$$N'_{i_1} \leq n^{1+\nu} \cdot n^{-(i_1-(i_0+1))\cdot\rho} \leq n^{\nu} \,, \tag{8}$$

the last phase i_1 contributes at most $N'_{i_1} \cdot N'_{i_1} \leq n^{2\nu} \leq n^{1+\nu}$ edges to the hopset. We conclude that the total number of edges is $O(n^{1+\nu})$.

Recall that the expected running time of the Dijkstra explorations at level $i < i_1$ is $O(|E| + n \log n)/p'_i$. Thus the expected running time of the first i_0 levels converges to $O(|E| + n \log n) \cdot n^\rho$, while each of the at most $\lceil 1/\rho \rceil + 1$ remaining levels will take also $O(|E| + n \log n) \cdot n^\rho$ time. The final level is expected to take $O(|E| + n \log n) \cdot n^\nu$ as well, since there are expected $O(n^\nu)$ vertices at A_{i_1} from which Dijkstra is performed. The price we pay is in a higher number of sets, which increases the exponent of β by at most an additive $1/\rho$. The following result summarizes the discussion.

Theorem 3. *For any weighted graph $G = (V, E)$ on n vertices, and any $1 < k \leq \log \log n - 2$, $0 < \rho < 1$, there is a randomized algorithm that runs in expected time $O(|E| + n \log n) \cdot n^\rho/\rho$, and computes an edge set H of size at most $O(n^{1+1/(2^k-1)})$. This edge set H is a (β, ϵ)-hopset for any $0 < \epsilon < 1$, where*

$$\beta = O\left(\frac{k + 1/\rho}{\epsilon}\right)^{k+1/\rho+1}.$$

By substituting $\kappa = 2^k - 1$, we obtain an improved version of the hopsets of [EN16a], where both the size of the hopset and the running time are smaller by a factor of $\log n$, while the other parameters remain the same. Another notable advantage is that it yields a single hopset which works for all $0 < \epsilon < 1$ simultaneously.

Corollary 3. *For any weighted graph $G = (V, E)$ on n vertices, and any $\kappa \geq 1$, $0 < \rho < 1$, there is a randomized algorithm that runs in expected time $O(|E| + n \log n) \cdot n^\rho/\rho$, and computes H of size at most $O(n^{1+1/\kappa})$, which is a (β, ϵ)-hopset for any $0 < \epsilon < 1$, where*

$$\beta = O\left(\frac{\log \kappa + 1/\rho}{\epsilon}\right)^{\log \kappa+1/\rho+1}.$$

3 Distributed Models

We will consider two standard models in distributed computing: the Congested Clique model, and the CONGEST model. In both models every vertex of an n-vertex graph $G = (V, E)$ hosts a processor, and the processors communicate with one another in discrete rounds, via short messages. Each message is allowed to contain an identity of a vertex, an edge weight, a distance in the graph, or anything else of no larger (up to a fixed constant factor) size.[6] The local computation is assumed to require zero time, and we are interested in algorithms that run for as few rounds as possible. In the Congested Clique model, we assume that all vertices are interconnected via direct edges, while in the CONGEST model, every vertex can send messages only to its G-neighbors (the weight of edges is irrelevant to the communication time).

3.1 Congested Clique Model

We first show how to construct the hopset in the Congested Clique model. In order to avoid a high number of rounds when computing distances for determining the bunches $\{B(u)\}$, we built the hopset in $\log n$ levels, where each level ℓ hopset will only "take care" of pairs that have a shortest path with *at most* 2^ℓ *hops*. This is somewhat different from previous works [Coh00, Nan14, HKN16, EN16a], in which the level ℓ hopset handled pairs with *distance in the range* $[2^\ell, 2^{\ell+1}]$. A few advantages of our current approach: it easily

[6] Typically, in the CONGEST model only messages of size $O(\log n)$ bits are allowed, but edge weights are restricted to be at most polynomial in n. Our definition is geared to capture a more general situation, when there is no restriction on the aspect ratio. Hence results achieved in our more general model are more general than previous ones.

avoids the dependency on the ratio between largest and smallest distance, and also the final hopset is just the level $\log n$ one, so we can obtain a linear size hopset (unlike previous works which took the union of all levels).

There are a few technical difficulties in implementing the algorithm of Section 2 in a distributed setting. The first is that the [TZ01a] method for computing the bunches $B(u)$ was to compute their "inverses" – called clusters.[7] Alas, it is not known how to compute these clusters in a distributed manner when errors are allowed. Rather, we compute the bunches directly, and to avoid the potential large congestion (a vertex may be a part of many bunches, and needs to send messages for all of them), we replace the bunches with *half-bunches* (i.e., taking only points closer than half the distance to the pivot). See below for the formal definition. The second issue (of congestion) is more subtle, and arises since hop-bounded distances do not obey the triangle inequality. For the stretch analysis to go through, we need that the weight of the hopset edges will be bounded by a certain path between the end-points of the edge (see (14)). In order to ensure that this happens, we build each hopset for level ℓ gradually, i.e. the bunches are created first for $A_0 \setminus A_1$, then for $A_1 \setminus A_2$ and so on, where each time the partial hopset is added to the graph on which we compute distance.

We say that H is a (β, ϵ, t)-*hopset* if $G \cup H$ provides $(1+\epsilon)$ approximation with at most β hops for all pairs $x, y \in V$ such that $d_G(x,y) = d_G^{(t)}(x,y)$ (i.e., the pairs that have a shortest path consisting of at most t edges between them). Note that the empty set is a $(1,0,1)$-hopset (and thus also a $(\beta,\epsilon,1)$-hopset for all $\beta \geq 1$ and $\epsilon > 0$). Given a $(\beta, \epsilon_{\ell-1}, 2^{\ell-1})$-hopset $H^{(\ell-1)}$, we build a $(\beta, \epsilon_\ell, 2^\ell)$-hopset $H^{(\ell)}$, where $1+\epsilon_\ell = (1+\epsilon)^\ell$ for some $0 < \epsilon < 1/(5\log n)$. The final hopset will be $H^{(\log n)}$. (We stress that the previous hopsets $H^{(1)}, \ldots, H^{(\ell-1)}$ are only used to compute $H^{(\ell)}$, and are not contained in it.)

Observe that $H^{(\ell-1)}$ is a $(2\beta, \epsilon_{\ell-1}, 2^\ell)$-hopset, since every path with at most 2^ℓ hops can be partitioned into two paths of at most $2^{\ell-1}$ hops each, and $H^{(\ell-1)}$ provides a $1+\epsilon_{\ell-1}$ approximation with β hops for each of these. It follows that for any $x, y \in V$,

$$d^{(2\beta)}_{G \cup H^{(\ell-1)}}(x,y) \leq (1+\epsilon_{\ell-1}) \cdot d_G^{(2^\ell)}(x,y) \; . \tag{9}$$

We sample sets $V = A_0 \supseteq A_1 \supseteq \cdots \supseteq A_{k'+1} = \emptyset$ as in Section 2.2, where $k' = i_1 \leq k + 1/\rho + 1$ is the number of sets. We introduce a subtle change to the construction – in the previous section we defined for each $u \in A_i \setminus A_{i+1}$ a set $B(u)$, and added the edges (u,v) for all $v \in B(u)$, simultaneously for all $0 \leq i \leq k'$. Here we shall build the hopset $H = H^{(\ell)}$ gradually: For each $i = 0, 1, \ldots, k'$ we define a set of edges $H_i = H_i^{(\ell)}$ corresponding to the bunches of vertices in $V \setminus A_{i+1}$, and finally take $H = H_{k'}$.

Fix some $0 \leq i \leq k'$, and assume we built the set H_{i-1} (when $i = 0$ define $H_{-1} = \emptyset$). We shall work in the graph G_i, defined by

$$G_i = G \cup H^{(\ell-1)} \cup H_{i-1} \; .$$

The algorithm consists of two stages. On the first stage, for 8β rounds, run a Bellman-Ford exploration in G_i rooted at A_{i+1}, to obtain for each $u \in V$ the value $\hat{d}(u, A_{i+1}) = d^{(8\beta)}_{G_i}(u, A_{i+1})$. (If some vertex $v \in V$ was not found in the exploration, then we set $\hat{d}(v, A_{i+1}) = \infty$.) Also for each vertex $u \in A_i \setminus A_{i+1}$ with $\hat{d}(u, A_{i+1}) < \infty$, store $p(u)$ as a vertex $p \in A_{i+1}$ satisfying $\hat{d}(u, A_{i+1}) = d^{(8\beta)}_{G_i}(u,p)$.

To determine the bunches (this is the second stage), each vertex $u \in A_i \setminus A_{i+1}$ conducts another Bellman-Ford exploration in the graph G_i rooted at u, this time for only 4β rounds, i.e., half the number of hops 8β of the first exploration, to distance less than $\hat{d}(u, A_{i+1})/2$ (i.e., the messages whose origin is u contain the

[7] The cluster $C(v)$ is defined as follows: each point $u \in C(v)$ iff $v \in B(u)$.

value $\hat{d}(u, A_{i+1})$, and only vertices within this distance from u forward the message in the next round). Define the half-bunch $B(u) = \{v \in A_i \ : \ d^{(4\beta)}_{G_i}(u, v) < \hat{d}(u, A_{i+1})/2\} \cup \{p(u)\}$. Let

$$H_i = H_{i-1} \cup \bigcup_{u \in A_i \setminus A_{i+1}} \{(u, v) \ : \ v \in B(u)\} \ , \tag{10}$$

and set the weight of the edge (u, v) as the distance discovered in the exploration (i.e., $d^{(4\beta)}_{G_i}(u, v)$ for $v \in B(u)$ and $d^{(8\beta)}_{G_i}(u, p(u))$ for the pivot).

Claim 4. $\mathbb{E}[|H|] \leq O(n^{1+\nu})$.

Proof. The argument is essentially the same as the one in Section 2.2. The only difference is that when analyzing in step $0 \leq i < k'$ the expected size of a bunch, i.e., $\mathbb{E}[|B(u)|]$ for some $u \in A_i \setminus A_{i+1}$, we consider the ordering on V given by the 4β-bounded distance from u in the graph G_i, i.e., according to $d^{(4\beta)}_{G_i}(u, \cdot)$. Then the size of $B(u)$ is bounded by the index of the first vertex in this ordering that is included in A_{i+1}. Since every $v \in A_i$ is included in A_{i+1} independently with probability p'_i, we have that $\mathbb{E}[|B(u)|] \leq 1/p'_i$. In fact, $B(u)$ may have a smaller size than the first index included in A_{i+1}, since we use more hops for computing distances to pivots (which reduces the distance threshold for being in $B(u)$), and since we only take into the bunch points that are less than half the distance to the pivot. Finally, for the last level k' we have that for $u \in A_{k'}$, $\mathbb{E}[|B(u)|] \leq \mathbb{E}[|A_{k'}|] \overset{(8)}{\leq} n^{\nu}$. Combining this with bounds on $N'_i = \mathbb{E}[|A_i|]$ in Section 2.2 we can bound the size of H in the same manner. ☐

In fact, since $|B(u)|$ is stochastically bounded by a geometric distribution with parameter $p'_i \geq n^{-\rho}$, it follows that with high probability for all $v \in V$,

$$|B(v)| \leq 4n^{\rho} \cdot \ln n \ . \tag{11}$$

Claim 5. *The number of rounds required is whp* $O(n^{\rho} \cdot k' \cdot \log^2 n \cdot \beta)$.

Proof. The sampling of the sets A_i is done independently for each vertex, therefore it requires no communication. For each $1 \leq \ell \leq \log n$ and $0 \leq i < k'$, we conduct a single Bellman-Ford exploration in G_i rooted at A_{i+1} for 8β rounds. Since in the Congested Clique model all edges are present, this requires $O(\beta)$ rounds per exploration (every vertex sends just a single message to all its neighbors every round). The more expensive step is the explorations to range 4β rooted at each $u \in A_i \setminus A_{i+1}$. The number of rounds in these explorations is affected by the number of messages a vertex $v \in V$ needs to forward to its neighbors at each round. In what follows we prove that with high probability this number is at most $O(n^{\rho} \log n)$ for every $v \in V$.

Fix $0 \leq i < k'$. Consider $v \in V$, and order the vertices of A_i according to their 4β-bounded distance to v in G_i, that is, according to $d^{(4\beta)}_{G_i}(v, \cdot)$. Since $p'_i \geq n^{-\rho}$, the probability that none of the first $2n^{\rho} \ln n$ vertices in that ordering is sampled to A_{i+1} is at most

$$\left(1 - n^{-\rho}\right)^{2n^{\rho} \ln n} \leq 1/n^2 \ .$$

So by the union bound on the n vertices, with high probability, for each $v \in V$ at least one of the first $2n^{\rho} \ln n$ vertices in its ordering of A_i is sampled to A_{i+1}. Denote by $z \in A_{i+1}$ the first vertex in the ordering of v that was chosen to A_{i+1}. We claim that no vertex $u \in A_i$, that appears after z in the ordering

of v, will cause v to forward messages concerning $B(u)$. This is because in the first stage we performed the Bellman-Ford exploration rooted at A_{i+1} for 8β rounds. Thus

$$\hat{d}(u, A_{i+1}) \leq d_{G_i}^{(8\beta)}(u, z) \leq d_{G_i}^{(4\beta)}(u, v) + d_{G_i}^{(4\beta)}(v, z) \leq 2d_{G_i}^{(4\beta)}(u, v) \ ,$$

where the last inequality uses the assumption that u appeared after z in v's ordering. We obtained $d_{G_i}^{(4\beta)}(u, v) \geq \hat{d}(u, A_{i+1})/2$. So by the definition of half bunch, $v \notin B(u)$ and thus, v will not forward u's messages.

We still have to argue about the last level $i = k'$ (since no vertex is chosen to $A_{k'+1}$). Recall that the expected size of $A_{k'}$ is bounded by n^{ν}, as shown in (8). It can be easily checked that whp $|A_{k'}| \leq O(n^{\nu} \cdot \log n) \leq O(n^{\rho} \cdot \log n)$. (Recall that $\rho \geq 2\nu$.) We conclude that whp, every vertex needs to send at most $O(n^{\rho} \log n)$ messages to implement a single step of Bellman-Ford. There are $O(\beta)$ rounds for each $\ell = 1, \ldots, \log n$ and each $0 \leq i \leq k'$, so the total number of rounds required is $O(n^{\rho} \cdot k' \cdot \log^2 n \cdot \beta)$. □

Next, we prove an analogue of Lemma 1 for the distributed setting. There are several subtle differences described in the beginning of this section. So we provide a complete proof that addresses these subtleties.

Lemma 6. *Fix any $0 < \delta < 1/(15k')$, set $\beta = (3/\delta)^{k'}$, and let $x, y \in V$ be such that $d_G(x, y) = d_G^{(2^\ell)}(x, y)$. Then for every $0 \leq i \leq k'$, at least one of the following two assertions holds:*

1. $d_{G \cup H_i}^{((3/\delta)^i)}(x, y) \leq (1 + \epsilon_{\ell-1}) \cdot (1 + 12\delta i) \cdot d_G(x, y)$.
2. *There exists $z \in A_{i+1}$ such that $d_{G \cup H_i}^{((3/\delta)^i)}(x, z) \leq 3d_G(x, y)$.*

Proof. The proof is by induction on i. We start with the base case $i = 0$. If it is the case that $x \in A_1$ then we can take $z = x$ and the second item holds trivially. Otherwise, consider the case that $x \in A_0 \setminus A_1$ and $y \in B(x)$. Then we added an edge (x, y) to H_0 of weight $d_{G_0}^{(4\beta)}(x, y)$ (the case that $y = p(x)$ is similar, replacing 4β by 8β). Recall that $G_0 = G \cup H^{(\ell-1)}$. Hence

$$d_{H_0}^{(1)}(x, y) = d_{G_0}^{(4\beta)}(x, y) \leq d_{G_0}^{(2\beta)}(x, y) \overset{(9)}{\leq} (1 + \epsilon_{\ell-1}) \cdot d_G^{(2^\ell)}(x, y) = (1 + \epsilon_{\ell-1}) \cdot d_G(x, y) \ ,$$

so the first item holds. The last case is that $x \in A_0 \setminus A_1$ and $y \notin B(x)$. By definition of H_0, it must be that

$$\hat{d}(x, A_1) \leq 2d_{G_0}^{(4\beta)}(x, y) \ . \tag{12}$$

Since $p(x) \in B(x)$, the edge $(x, p(x))$ is in the hopset, and its weight is

$$d_{G_0}^{(8\beta)}(x, p(x)) = \hat{d}(x, A_1) \overset{(12)}{\leq} 2d_{G_0}^{(4\beta)}(x, y) \leq 2(1 + \epsilon_{\ell-1}) \cdot d_G(x, y) < 3d_G(x, y) \ ,$$

where for the last two inequalities we again used (9) and the fact that $1 + \epsilon_{\ell-1} < 3/2$. (The latter holds since we assume $\epsilon < 1/(5 \log n)$ and $\ell \leq \log n$, so $(1 + \epsilon_\ell) = (1 + \epsilon)^\ell < e^{1/5}$). This proves the second item with $z = p(x) \in A_1$.

Assume the claim holds for i, and we prove for $i + 1$. Consider the shortest path $\pi(x, y)$ between x, y in G that contains at most 2^ℓ edges, and partition it into $J \leq 1/\delta$ segments $\{L_j = [u_j, v_j]\}_{j \in [J]}$ as in the proof of Lemma 1. We use the induction hypothesis for all pairs (u_j, v_j) with parameter i. (By the virtue of lying on a shortest path that has at most 2^ℓ edges, all these pairs satisfy $d_G^{(2^\ell)}(u_j, v_j) = d_G(u_j, v_j)$). Consider first the case that the first item holds for all of them, that is, $d_{G \cup H_i}^{((3/\delta)^i)}(u_j, v_j) \leq (1 + \epsilon_{\ell-1}) \cdot (1 + 12\delta i) \cdot d_G(u_j, v_j)$.

Then we take the path in $G \cup H_i$ that consists of the $(3/\delta)^i$-hops between each pair u_j, v_j, and the edges (v_j, u_{j+1}) of G. Since by (10), $H_i \subseteq H_{i+1}$, we have

$$d^{((3/\delta)^{i+1})}_{G \cup H_{i+1}}(x,y) \leq \sum_{j \in [J]} (d^{((3/\delta)^i)}_{G \cup H_i}(u_j, v_j) + d^{(1)}_G(v_j, u_{j+1})) \leq (1+\epsilon_{\ell-1}) \cdot (1+12\delta i) \cdot d_G(x,y) \ ,$$

which concludes the proof for the first case. The second case is that there are pairs (u_j, v_j) for which only the second item holds. Let $l \in [J]$ (resp., $r \in [J]$) be the first (resp., last) index for which the first item does not hold for the pair (u_l, v_l) (resp., (u_r, v_r)). Then there are $z_l, z_r \in A_{i+1}$ such that

$$d^{((3/\delta)^i)}_{G \cup H_i}(u_l, z_l) \leq 3 d_G(u_l, v_l) \ \text{ and } \ d^{((3/\delta)^i)}_{G \cup H_i}(v_r, z_r) \leq 3 d_G(u_r, v_r) \ . \tag{13}$$

Consider first the case that $z_l \in A_{i+2}$. Then we take $z = z_l$, and derive

$$\begin{aligned} d^{((3/\delta)^{i+1})}_{G \cup H_{i+1}}(x,z) &\leq \sum_{j=1}^{l-1} \left(d^{((3/\delta)^i)}_{G \cup H_i}(u_j, v_j) + d^{(1)}_G(v_j, u_{j+1}) \right) + d^{((3/\delta)^i)}_{G \cup H_i}(u_l, z_l) \\ &\overset{(13)}{\leq} (1+\epsilon_{\ell-1}) \cdot (1+12\delta i) \cdot d_G(x, u_l) + 3 d_G(u_l, v_l) \\ &\leq 3 d_G(x,y) \ , \end{aligned}$$

where in the second inequality we used that the first item holds for all intervals until the l-th one, and in the final one that $1 + \epsilon_{\ell-1} < 3/2$ and $1 + 12\delta i < 2$.

From now on assume $z_l \in A_{i+1} \setminus A_{i+2}$. Recall that the Bellman-Ford explorations that constructed H_{i+1} were conducted in the graph $G_{i+1} = G \cup H^{(\ell-1)} \cup H_i$. These explorations were conducted to hop-depth 8β on the first stage, and 4β on the second. This allows us to provide the following bound:

$$\begin{aligned} d^{(4\beta)}_{G_{i+1}}(z_l, z_r) &\leq d^{(\beta)}_{G \cup H_i}(z_l, u_l) + d^{(2\beta)}_{G \cup H^{(\ell-1)}}(u_l, v_r) + d^{(\beta)}_{G \cup H_i}(v_r, z_r) \\ &\overset{(9)}{\leq} d^{((3/\delta)^i)}_{G \cup H_i}(z_l, u_l) + (1+\epsilon_{\ell-1}) \cdot d_G(u_l, v_r) + d^{((3/\delta)^i)}_{G \cup H_i}(v_r, z_r) \ . \end{aligned} \tag{14}$$

Here the first inequality follows by the triangle inequality, the second uses that $(3/\delta)^i \leq \beta$, that u_l, v_r lie on a shortest path with at most 2^ℓ hops, and that $H^{(\ell-1)}$ is a $(2\beta, \epsilon_{\ell-1}, 2^\ell)$ hopset.

Consider the case that $z_r \in B(z_l)$, then we have a hopset edge (z_l, z_r) that was introduced in H_{i+1}. In particular, since we used 4β steps in the exploration from z_l, we have that

$$d^{(1)}_{H_{i+1}}(z_l, z_r) = d^{(4\beta)}_{G_{i+1}}(z_l, z_r) \overset{(14)}{\leq} d^{((3/\delta)^i)}_{G \cup H_i}(z_l, u_l) + (1+\epsilon_{\ell-1}) \cdot d_G(u_l, v_r) + d^{((3/\delta)^i)}_{G \cup H_i}(v_r, z_r) \ . \tag{15}$$

Next, apply the inductive hypothesis on segments $\{L_j\}$ for $j < l$ and $j > r$, and in between use the detour via u_l, z_l, z_r, v_r. Since there are at most $1/\delta - 1$ intervals for which we use the first item in the inductive hypothesis, the total number of hops we will need is at most $(1/\delta - 1) \cdot (3/\delta)^i + 1/\delta + 2(3/\delta)^i + 1$.

This is at most $(3/\delta)^{i+1}$ whenever $\delta < 1/2$. It follows that

$$\begin{aligned}
d^{((3/\delta)^{i+1})}_{G\cup H_{i+1}}(x,y) &\leq \sum_{j=1}^{l-1}\left[d^{((3/\delta)^i)}_{G\cup H_i}(u_j,v_j) + d^{(1)}_G(v_j,u_{j+1})\right] + d^{((3/\delta)^i)}_{G\cup H_i}(u_l,z_l) + d^{(1)}_{H_{i+1}}(z_l,z_r) \\
&\quad + d^{((3/\delta)^i)}_{G\cup H_i}(z_r,v_r) + d^{(1)}_G(v_r,u_{r+1}) + \sum_{j=r+1}^{J}\left[d^{((3/\delta)^i)}_{G\cup H_i}(u_j,v_j) + d^{(1)}_G(v_j,u_{j+1})\right] \\
&\overset{(15)}{\leq} (1+\epsilon_{\ell-1})\cdot\left[(1+12\delta i)\cdot d_G(x,u_l) + d_G(u_l,v_r) + (1+12\delta i)\cdot d_G(v_r,y)\right] \\
&\quad + 2d^{((3/\delta)^i)}_{G\cup H_i}(u_l,z_l) + 2d^{((3/\delta)^i)}_{G\cup H_i}(z_r,v_r) \\
&\overset{(13)}{\leq} (1+\epsilon_{\ell-1})\cdot(1+12\delta i)\cdot d_G(x,y) + 12\delta\cdot d_G(x,y) \\
&\leq (1+\epsilon_{\ell-1})\cdot(1+12\delta(i+1))\cdot d_G(x,y)\ .
\end{aligned}$$

In the penultimate inequality we used that both $d_G(u_l,v_l), d_G(u_r,v_r) \leq \delta\cdot d_G(x,y)$. This demonstrates that item 1 holds in this case.

The final case to consider is that $z_r \notin B(z_l)$ (and $z_l \in A_{i+1}\setminus A_{i+2}$). Let $z = p(z_l) \in A_{i+2}$. Since $z_r \in A_{i+1}$, the definition of $B(z_l)$ implies that

$$d^{(1)}_{H_{i+1}}(z_l,z) = \hat{d}(z_l,A_{i+2}) \leq 2d^{(4\beta)}_{G_{i+1}}(z_l,z_r)\ . \tag{16}$$

(Recall that $G_{i+1} = G\cup H^{(\ell-1)}\cup H_i$.)

We now claim that item 2 holds for such a choice of z. Indeed, by (13), we have

$$3\cdot d^{(3/\delta)^i}_{G\cup H_i}(u_l,z_l) + 2\cdot d^{(3/\delta)^i}_{G\cup H_i}(v_r,z_r) \ \leq\ 15\cdot d_G(x,y)\ . \tag{17}$$

Hence,

$$\begin{aligned}
d^{((3/\delta)^{i+1})}_{G\cup H_{i+1}}(x,z) &\leq \sum_{j=1}^{l-1}\left[d^{((3/\delta)^i)}_{G\cup H_i}(u_j,v_j) + d^{(1)}_G(v_j,u_{j+1})\right] + d^{((3/\delta)^i)}_{G\cup H_i}(u_l,z_l) + d^{(1)}_{H_{i+1}}(z_l,z) \\
&\overset{(16)\wedge(14)}{\leq} (1+\epsilon_{\ell-1})\cdot(1+12\delta i)\cdot d_G(x,u_l) + d^{((3/\delta)^i)}_{G\cup H_i}(u_l,z_l) \\
&\quad + 2\cdot\left[d^{((3/\delta)^i)}_{G\cup H_i}(z_l,u_l) + (1+\epsilon_{\ell-1})\cdot d_G(u_l,v_r) + d^{((3/\delta)^i)}_{G\cup H_i}(v_r,z_r)\right] \\
&\overset{(17)}{\leq} (1+\epsilon_{\ell-1})\cdot(1+12\delta i)\cdot d_G(x,u_l) + 15\delta\cdot d_G(x,y) + 2(1+\epsilon_{\ell-1})\cdot d_G(u_l,v_r) \\
&\leq 3d_G(x,y)\ ,
\end{aligned}$$

where the last inequality we used that $\delta < 1/(15k)$, $k \geq 2$ and $1+\epsilon_{\ell-1} < e^{1/5} < 5/4$, so that both $(1+\epsilon_{\ell-1})\cdot(1+12\delta i) + 15\delta \leq 3$ and $2(1+\epsilon_{\ell-1}) + 15\delta \leq 3$. □

Taking $\delta = \epsilon/(15k')$ and picking $i = k'$, the second item of Lemma 6 cannot hold for any $x,y\in V$ (because $A_{k'+1} = \emptyset$), so we have for every $x,y\in V$ such that $d^{(2^\ell)}_G(x,y) = d_G(x,y)$ that

$$d^{(\beta)}_{G\cup H^{(\ell)}}(x,y) \leq (1+\epsilon_{\ell-1})\cdot(1+\epsilon)\cdot d^{(2^\ell)}_G(x,y) = (1+\epsilon_\ell)\cdot d_G(x,y)\ .$$

Recall that $\beta = (3/\delta)^{k'} = (45k'/\epsilon)^{k'}$. Rescaling $\epsilon' = \epsilon/\log n$ and taking $\ell = \log n$, we derive the following theorem.

Theorem 4. *For any weighted graph $G = (V, E)$ on n vertices, an integer $k > 1$, and parameters $0 < \rho < 1$, $0 < \epsilon < 1/5$, there is a distributed algorithm in the Congested Clique model running in $\tilde{O}(n^{\rho} \cdot \beta)$ rounds, that computes a (β, ϵ)-hopset H of size at most $O(n^{1+1/(2^k-1)})$, where*

$$\beta = O\left(\frac{(k + 1/\rho) \cdot \log n}{\epsilon}\right)^{k+1/\rho+1}. \tag{18}$$

Remark 1. *We note that by (11) and Claim 5, the memory requirement from every vertex is $\tilde{O}(n^{\rho})$. This is because the latter shows that this is a bound on the number of messages every vertex needs to send in each round, and the former indicates that whp storing $B(v)$ for any $v \in V$ requires only so much space.*

We remark that one can achieve β independent of n by either applying the construction recursively, as we do in Section 4 for the parallel implementation, or by using an idea from [EN16a]. We next describe the latter: fix a parameter t, and use the hopset $H^{(\ell)}$ to compute the hopset $H^{(\ell+t)}$; Since $H^{(\ell)}$ is also a $(2^t \cdot \beta, \epsilon, 2^{\ell+t})$-hopset, we need explorations to range $2^t \cdot \beta$ in order for an appropriate variant of (9) to hold. There will be only $(\log n)/t$ levels until $H^{(\log n)}$ is built, so we gain a factor of t in β. We derive the following result.

Theorem 5. *For any weighted graph $G = (V, E)$ on n vertices, integers $k > 1$, $t \geq 1$, and parameters $0 < \rho < 1$, $0 < \epsilon < 1/5$, there is a distributed algorithm in the Congested Clique model that runs in $\tilde{O}(n^{\rho} \cdot \beta \cdot 2^t/t)$ rounds, and computes H of size at most $O(n^{1+1/(2^k-1)})$, which is a (β, ϵ)-hopset, where*

$$\beta = O\left(\frac{(k + 1/\rho) \cdot \log n}{t \cdot \epsilon}\right)^{k+1/\rho}.$$

In particular, taking $t = \rho \log n$ and rescaling $\rho' = 2\rho$, gives

Corollary 7. *For any weighted graph $G = (V, E)$ on n vertices, an integer $k > 1$, and parameters $0 < \rho < 1/2$, $0 < \epsilon < 1/5$, there is a distributed algorithm in the Congested Clique model that runs in $\tilde{O}(n^{\rho} \cdot \beta)$ rounds, and computes H of size at most $O(n^{1+1/(2^k-1)})$, which is a (β, ϵ)-hopset, where*

$$\beta = O\left(\frac{k + 1/\rho}{\rho \cdot \epsilon}\right)^{k+2/\rho}.$$

3.2 CONGEST Model

Given a weighted graph $G = (V, E, w)$ representing the network, in the CONGEST model we will be interested in a setting where there is a "virtual" graph $G' = (V', E', w')$ embedded in G, i.e., $V' \subseteq V$. We would like to construct a hopset for G'. It is motivated by distributed applications of hopsets for approximate shortest paths computation, distance estimation and routing [HKN14, Nan14, HKN16, LP15, EN16b, EN16a], which require a hopset for a virtual graph embedded in the underlying network in the above way.

In a similar manner to [EN16a], we can modify our algorithm in the Congested Clique model to the CONGEST model. The following lemma provides a way to perform Bellman-Ford exploration using small memory.

Lemma 8. *Let $G'' = (V', E' \cup H)$ be a virtual graph on m vertices embedded in a graph $G = (V, E)$ of hop-diameter D, such that edges in E' correspond to B-bounded distances in G, and H has arboricity α (i.e., one can orient the edges of H to have out-degree at most α). Moreover, every vertex $v' \in V'$ knows at most α its outgoing edges in H. Then one can compute β iterations of Bellman-Ford in G'' in the CONGEST model within $O(m \cdot \alpha + B + D) \cdot \beta \cdot \log n$ rounds, so that every vertex requires only $O(\alpha \log n)$ memory.*

Proof. To implement a single iteration of the Bellman-Ford exploration, every vertex $v \in V'$, which holds a current distance estimate, will need to communicate it to its neighbors in G''. First, it will initiate an exploration in G for B rounds. In each round, every vertex $u \in V$ will forward the smallest value it received so far. This guarantees that if $\{v, w\} \in E'$, then w will receive v's message (or a smaller value).

We now have to handle the edges of H. Let T be a spanning tree of G with hop-depth D. Every $v \in V'$ will broadcast via T its value to the entire graph, and will also send all the existing edges of H incident on it that v knows about. All vertices $w \in V'$ that know of a hopset edge $\{v, w\}$ (or that learn about it from v's message) will update their value accordingly. Since there are $O(m \cdot \alpha)$ messages, this can be done in $O(m \cdot \alpha + D)$ rounds. In order to guarantee small internal memory, each v selects at random a number from $\{1, 2, \ldots, m \cdot \alpha\}$ for each message it sends, as a round to start its broadcast (clearly this increases the number of rounds by at most $m \cdot \alpha$). Since each message of v will reach every vertex of T at most once, the probability that some $u \in V$ receives t messages in a single round is at most $\binom{m \cdot \alpha}{t} \cdot 1/(m \cdot \alpha)^t \leq (e/t)^t$. Thus, with high probability, no vertex will receive more than $O(\log n)$ messages each round. By increasing the number of rounds by $O(\log n)$, whp there will be no congestion. The total number of rounds required is thus $O(m \cdot \alpha + B + D) \cdot \beta \cdot \log n$. □

We now show how to use Lemma 8 to construct a hopset for G', in the setting where E' are edges corresponding to $B = \tilde{O}(m)$-bounded distances in G (without computing G' explicitly). Recall that in the ith iteration of constructing $H = H^{(\ell)}$, we have already built the previous hopset $H^{(\ell-1)}$ and the partial hopset H_{i-1}. Since we desire limited memory, every vertex v stores only the "outgoing" hopset edges, those to vertices in its bunch $B(v)$. Recall that by (11), whp $|B(v)| \leq O(m^\rho \cdot \log n)$, for all $v \in V'$.

We work in the graph $G_i = G' \cup H^{(\ell-1)} \cup H_{i-1}$. In order to implement the $O(\beta)$-bounded exploration rooted at A_{i+1} (the second stage of the ith iteration), we simply apply Lemma 8 on G_i with $\alpha = O(m^\rho \cdot \log n)$. The explorations from vertices of $A_i \setminus A_{i+1}$ (the second stage of the ith iteration) are done in a similar manner. However, there is a larger congestion than in the first stage, due to the multiple sources of limited explorations. Recall that in the limited exploration whose origin is $v \in A_i \setminus A_{i+1}$, each intermediate node $x \in V'$ forwards the message iff its current estimate is strictly less than $\hat{d}(v, A_{i+1})/2$ (this value is part of the message v sends). We enforce the exact same rule for vertices $u \in V$ as well. If a message concerning v should pass in G' from x to its neighbor y, then all vertices on the B-bounded path in G that implements the edge $(x, y) \in E'$ will have estimates smaller than that of y, therefore will forward the message on. In the proof of Claim 5 we saw that each $x \in V'$ participates whp in at most $O(m^\rho \cdot \log n)$ explorations for each iteration i. The argument is identical for $u \in V$ as well, so the congestion induced in the first stage of Lemma 8 (the exploration in G for B rounds) by multiple sources is only $O(m^\rho \cdot \log n)$. Note that in the second phase (broadcasting the edges of H), the number of messages increases to $O(m \cdot \alpha + m \cdot m^\rho \cdot \log n)$. Thus, the total number of rounds required is still $\tilde{O}(m^{1+\rho} + D) \cdot \beta$. We summarize the discussion with the following result.

Theorem 6. *For any weighted graph $G = (V, E)$ with hop-diameter D, an integer $k > 1$, and parameters $0 < \rho < 1$, $0 < \epsilon < 1/5$, and (an implicit) virtual graph $G' = (V', E')$ embedded in G on $|V'| = m$ vertices, there is a distributed algorithm in the CONGEST model that runs in $\tilde{O}(m^{1+\rho} + D) \cdot \beta$ rounds, that computes H, which is a (β, ϵ)-hopset for G', of size at most $O(m^{1+1/(2^k-1)})$, where*

$$\beta = O\left(\frac{(k + 1/\rho) \cdot \log m}{\epsilon}\right)^{k+1/\rho+1}.$$

Remark 2. *In the case that E' corresponds to $B = \tilde{O}(m)$-bounded distances in G, the hopset can be computed where every vertex has internal memory $\tilde{O}(m^\rho)$.*

Path-reporting Hopsets: Every hopset edge is implemented via some path in G. For our application to routing, we would like that every vertex on a path implementing a certain hopset edge will be aware of this hopset edge. This means that for every hopset edge $(x,y) \in H$, there exists a path P in G of length $w_H(x,y)$, and every vertex $u \in P$ knows about the hopset edge, and the distances $d_P(u,x)$, $d_P(u,y)$, and its neighbors on P. It was shown in [EN16a] how to adapt the Bellman-Ford exploration, so that paths information can be stored as well, at a cost of increasing the size of messages by a factor of $O(\beta)$. However, there was no guarantee on the number of hopset edges a vertex $u \in V$ can be a part of, which can be devastating when one desires small memory per vertex. We describe now an approach that eliminates the need for the message's size increase, and also ensures that each vertex belongs to a bounded number of paths that implement hopset edges. The issue that may cause a vertex to be in a path for many hopset edges, is that we use previous hopsets to construct a new one. Then the vertices implementing paths in these previous hopsets may not be discovered by the current explorations. So the argument of Claim 5 bounding the number of explorations that visit a certain vertex does not apply as is.

In order to guarantee that every $u \in V$ will need to store information for only $\tilde{O}(m^\rho)$ hopset edges, we need to slightly change the construction. First, we will define $H^{(\ell)} = H_{k'} \cup H^{(\ell-1)}$, so that every hopset will contain all the previous hopsets. (Recall that in our algorithm in Section 3.1 that computes non-path-reporting hopsets, we only used lower-scale hopsets to compute a higher-scale one. Once the mission of lower-scale hopsets was completed, they were ruthlessly erased.) Second, rather than performing the exploration from A_{i+1} in 8β steps, we apply Lemma 8 with $8\beta \cdot k' \cdot \log n + 1$ steps of Bellman-Ford. Note that in the proof of correctness we only used that there are at least 8β steps. Using more steps will only increase the number of rounds (by a poly-logarithmic factor). Recall that when computing the hopset $H^{(\ell)}$ at phase i, we have already computed H_{i-1}, and work in the graph $G_i = G' \cup H^{(\ell-1)} \cup H_{i-1}$. We can now argue that whp, there will not be too many hopset edges whose path in G contains u. The intuition is that the exploration from A_{i+1} has sufficiently many hops in order to discover this u, and so an argument similar to the one of Claim 5 will apply.

Fix $u \in V$, and order the vertices of A_i in increasing order according to their distance to u, where the distance from $v \in A_i$ to u is the shortest path consisting of at most $4\beta \cdot k' \cdot \log n$ edges of G_i and then at most B edges of G. Let z be the first vertex in that order that is included in A_{i+1}. We claim that the vertex u cannot belong to a path P that implements a hopset edge (x,y), such that $x \in A_i \setminus A_{i+1}$ is after z in the ordering of u.

Consider how the path P is built. One can initially start with $Q = \{(x,y)\}$, and then recursively replace the hopset edge in Q that contains u, with the 4β-bounded path in some G'_j that induces it. Note that this recursion depth is at most $k' \log n$, thus Q has at most $4\beta \cdot k' \log n$ edges. Since G_i contains all the edges of all previous hopsets, the exploration from A_{i+1} starting at z for $8\beta \cdot k' \cdot \log n + 1$ steps would have reached u after $4\beta \cdot k' \cdot \log n + 1$ edges of G_i (the B edges of G are an edge of G', and thus of G_i as well), and then after additional $4\beta \cdot k' \cdot \log n$ edges of G_i, it would have surely have reached y (because z is closer to u than x). We conclude that

$$\hat{d}(y, A_{i+1}) \leq d_{G_i}^{(8\beta \cdot k' \cdot \log n + 1)}(y,z) \leq d_P(y,u) + \hat{d}(u, A_{i+1}) \leq d_P(y,x) + d_P(u,x) \leq 2d_{G_i}^{(4\beta)}(x,y) \ ,$$

which is a contradiction to the fact that y joins $B(x)$.

Next, we have to show that each u will indeed learn the relevant information on all hopset edges it implements. Assume inductively that for any hopset edge $(x,y) \in H^{(\ell-1)} \cup H_{i-1}$, if P is the path in G that implements this edge, then every $u \in P$ knows about the edge, $d_P(u,x)$, $d_P(u,y)$, and its neighbors on P. A new hopset edge $(x,y) \in H_i$ is created whenever the exploration rooted at $x \in A_i \setminus A_{i+1}$ discovers a vertex $y \in A_i$. Recall that this exploration is done in G_i for 4β rounds. Whenever y joins $B(x)$ it will

send an acknowledgement on the 4β-bounded path back to x in G_i (every vertex discovered by x takes note of its "parent", the vertex who sent it the message of x). The acknowledgement phase can take place after the exploration concludes, and it will induce congestion that is no larger that the congestion created when sending the messages, so the number of rounds will at most double. Now, every vertex v in the 4β-bounded path from y to x that receives y's acknowledgement, knows that the edge to its parent v' is part of the path implementing the hopset edge (x, y). Recall that the edge (v, v') is either an edge of G', which is discovered via a B-round exploration in G – in which case all vertices along the path in G from v to v' can update the relevant information about (x, y) when v does a B-round exploration in G (this is the acknowledgement step), or otherwise $(v, v') \in H^{(\ell-1)} \cup H_{i-1}$. In the latter case, v will broadcast that the edge (v, v') implements (x, y), and its distances to x, y. By the induction hypothesis, each vertex u' that implements a path P' for the hopset edge (v, v') knows about it and its distances to v, v', thus when u' hears this broadcast (which is sent to all vertices of V), it knows it implements P, and can computes distances to x and y.

We conclude that whp every vertex needs to store only the $\tilde{O}(m^\rho)$ hopset edges that it implements. Note that the final hopset $H^{(\log n)}$ can omit all the previous hopsets (which were used only for calculations). We summarize this discussion with the following theorem.

Theorem 7. *For any weighted graph $G = (V, E)$ with hop-diameter D, an integer $k > 1$, and parameters $0 < \rho < 1$, $0 < \epsilon < 1/5$, and (an implicit) virtual graph $G' = (V', E')$ embedded in G on $|V'| = m$ vertices, there is a distributed algorithm in the CONGEST model that runs in $\tilde{O}(m^{1+\rho} + D) \cdot \beta$ rounds, that computes H, which is a (β, ϵ)* path-reporting *hopset for G', of size at most $O(m^{1+1/(2^k-1)})$, where*

$$\beta = O\left(\frac{(k + 1/\rho) \cdot \log m}{\epsilon}\right)^{k+1/\rho+1}.$$

In the case that E' corresponds to $B = \tilde{O}(m)$-bounded distances in G, the hopset can be computed where every vertex has internal memory $\tilde{O}(m^\rho)$.

4 PRAM Model

The algorithm described in Section 3.1 can be easily adapted to the PRAM model. For each $\ell = 1, 2, \ldots, \log n$, we build the hopset $H^{(\ell)}$ based on the previous hopset $H^{(\ell-1)}$. Each of the $O(\beta)$-bounded Bellman-Ford explorations for constructing H_i can be implemented in parallel in $O(\beta)$ rounds, where the congestion of $\tilde{O}(n^\rho)$ per vertex translates to extra work (rather than multiplying the number of rounds, as was the case in distributed models). Since there are $\log n$ values of ℓ, and $k' \leq k + 1/\rho + 1$ steps in each level, the number of rounds is only $O((k + 1/\rho) \cdot \log n \cdot \beta)$. We have the following result.

Theorem 8. *For any weighted graph $G = (V, E)$ on n vertices, an integer $k > 1$, and parameters $0 < \rho < 1$, $0 < \epsilon < 1/5$, there is a parallel algorithm running in $O((k + 1/\rho) \cdot \log n \cdot \beta)$ rounds and has $\tilde{O}(|E| \cdot n^\rho)$ work, that computes H of size at most $O(n^{1+1/(2^k-1)})$, which is a (β, ϵ)-hopset, where*

$$\beta = O\left(\frac{(k + 1/\rho) \cdot \log n}{\epsilon}\right)^{k+1/\rho+1}. \tag{19}$$

We can also apply the construction recursively: If $H(1)$ is the hopset given by Theorem 8 with $\beta_1 = \beta$ given in (19), then apply the construction on the graph $G \cup H(1)$, but only for levels ℓ up to $\ell_2 = \log \beta_1$, to

obtain a hopset $H(2)$. Since for any $x, y \in V$ we have $d^{(\beta_1)}_{G\cup H(1)}(x,y) \leq (1+\epsilon)d_G(x,y)$, then adding both $H(1)$ and $H(2)$ guarantees $d^{(\beta_2)}_{G\cup H(1)\cup H(2)}(x,y) \leq (1+\epsilon)^2 d_G(x,y)$, where $\beta_2 = \left(\frac{3c\cdot(k+1/\rho)\cdot\log\beta_1}{\epsilon}\right)^{k+1/\rho+1}$, where c is the constant hidden by the $O(\cdot)$ notation in (19). This bound follows because ϵ needs to be rescaled by $3\ell_2 = 3\log\beta_1$; the rescaling by $\log\beta_1$ is to compensate for the number of levels, and by 3 to reduce the error from $(1+\epsilon)^2$ back to $1+\epsilon$. Continuing in this manner for the next level with $\ell_3 = \log\beta_2$ levels, we obtain in general a recursion for $\beta_{i+1} = \left(\frac{3c\cdot(k+1/\rho)\cdot\log\beta_i}{\epsilon}\right)^{k+1/\rho+1}$, and it can be shown by induction that as long as $log^{(i)}n \geq 3c\log(k+1/\rho)$ we have

$$\beta_i \leq \left(\frac{8c\cdot(k+1/\rho)^2\cdot\left[\log(3c(k+1/\rho)/\epsilon)+\log^{(i)}n\right]}{\epsilon}\right)^{k+1/\rho+1}.$$

After at most $t = \log^* n$ iterations, we get that $\beta_t = O\left(\frac{(k+1/\rho)^2}{\epsilon}\right)^{(1+o(1))\cdot(k+1/\rho)}$. To summarize, this yields a hopset with constant parameter β that is computed in $\mathrm{polylog}(n)$ rounds.

Theorem 9. *For any weighted graph $G = (V,E)$ on n vertices, an integer $k > 1$, and parameters $0 < \rho < 1$, $0 < \epsilon < 1/5$, there is a parallel algorithm running in $O(\left(\frac{(k+1/\rho)\cdot\log n}{\epsilon}\right)^{k+1/\rho+2})$ rounds and has $\tilde{O}(|E|\cdot n^\rho)$ work, that computes H of size at most $O(n^{1+1/(2^k-1)}\cdot\log^* n)$, which is a (β,ϵ)-hopset, where*

$$\beta = O\left(\frac{(k+1/\rho)^2}{\epsilon}\right)^{(1+o(1))\cdot(k+1/\rho)}. \tag{20}$$

5 Distributed Routing with Small Memory

Here we improve the results of [EN16b, LPP16], and devise a compact routing scheme that can be efficiently implemented in a distributed network. The previous result of [EN16b] provides, for any parameter k, a scheme with stretch $4k-5+o(1)$, labels of size $O(k\log^2 n)$ and routing tables of size $O(n^{1/k}\log^2 n)$. The computation time of this scheme is $(n^{1/2+1/k}+D)\cdot\min\{(\log n)^{O(k)}, 2^{\tilde{O}(\sqrt{\log n})}\}$ rounds (in the CONGEST model). One drawback of this result (and also of [LPP16], which obtained slightly weaker results), is that although the final memory requirement from each vertex is $\tilde{O}(n^{1/k})$, the preprocessing step requires high memory (at least $\Omega(\sqrt{n})$). Indeed, some of the classical works on compact routing schemes [ABNLP90] addressed the issue of each vertex having only a limited memory throughout the construction of the routing scheme (albeit their round complexity was at least linear in n). Here we present a distributed construction that has that desirable property, and in addition we improve both the label and table size by a logarithmic factor, almost matching the best known bounds of [TZ01a, Che13] that are computed in a sequential manner.

We briefly sketch the approach of [EN16b], and the current improvement allowing low memory and improved bounds. First, construct the Thorup-Zwick hierarchy $V = A_0 \supseteq A_1 \supseteq A_k = \emptyset$, where each vertex in A_{i-1} is sampled to A_i independently with probability $n^{-1/k}$. Then the cluster $C(v) = \{u \in V : d_G(u,v) < d_G(u,A_{i+1})\}$ for $v \in A_i \setminus A_{i+1}$ can be viewed as tree rooted at v. Computing this cluster is done by a limited Dijkstra exploration from v, i.e., only vertices in $C(v)$ continue the exploration of v. Routing from x to y is done by finding an appropriate cluster $C(v)$ containing both x, y, and routing in that tree. Whenever $i < k/2$, these trees have whp depth $\tilde{O}(\sqrt{n})$. Hence they can be easily computed in a distributed manner within $\tilde{O}(n^{1/2+1/k})$ rounds. The main issue is computing the clusters for $i \geq k/2$.

The method of [EN16b] was to work with a virtual graph G', whose vertices are $V' = A_{k/2}$, and whose edges correspond to $B = c \cdot \sqrt{n} \log n$-bounded distances in G between the vertices of V'. Then a hopset is computed for this virtual graph, which enables the computation of Bellman-Ford explorations in only $O(\beta)$ rounds. The fact that β-bounded distances can suffer $1+\epsilon$ stretch creates additional complications; one needs to define *approximate clusters*, and make sure that these approximate clusters correspond to actual trees in G. Finally, since the trees corresponding to $C(v)$ for the high level vertices $v \in A_i$, $i \geq k/2$, can have large depth, one needs to adapt the Thorup-Zwick routing scheme for trees [TZ01b]. In both [EN16b, LPP16] this adaptation induced a logarithmic factor to both the table and the label size.

Our improved result has two main ingredients. First, we do not explicitly construct G'; In both [EN16b, LPP16], computing the weights of edges in G' was a rather expensive step, and required large memory and induced a factor depending logarithmically on the aspect ratio to the running time. In addition, only approximate values were obtained. We observe that not all the edges of G' are required for the algorithm, and thus we do not compute G' at all. Rather we compute only those edges of G' that are really needed for either the hopset or for the routing hierarchy. (This idea is reminiscent of [Elk17], where the virtual graph is also never entirely computed.)

Instead, we conduct the explorations in G' by implementing in each iteration a B-bounded search in G, which not only saves memory and running time, but also simplifies the analysis, since now there is no error in the edge weights of G'. Second, our new tree-routing scheme has both improved label and routing table size, and can be computed with small memory. (For more details, see Section 5.1.) Our result is summarized below.

Theorem 10. *Let $G = (V, E)$ be a weighted graph with n vertices and hop-diameter D, and let $k > 1$ be a parameter. Then there exists a routing scheme with stretch at most $4k - 5 + o(1)$, labels of size $O(k \log n)$ and routing tables of size $O(n^{1/k} \log n)$, that can be computed in a distributed manner within $(n^{1/2+1/k} + D) \cdot (\log n)^{O(k)}$ rounds, such that every vertex has memory of size $\tilde{O}(n^{1/k})$.*

Alternatively, whenever $k \geq \sqrt{\log n / \log\log n}$, the number of rounds can be made $(n^{1/2+1/k} + D) \cdot 2^{\tilde{O}(\sqrt{\log n})}$ with memory $2^{\tilde{O}(\sqrt{\log n})}$ at each vertex.

In particular, taking $k = \delta \log n / \log\log n$ for a small constant δ yields $(n^{1/2+1/k} + D) \cdot n^{O(\delta)}$ rounds with $\mathrm{polylog}(n)$ memory per vertex.

Construction of Routing Scheme. Let $G = (V, E)$ be a weighted graph, fix $k > 1$. Sample a collection of sets $V = A_0 \supseteq A_1 \cdots \supseteq A_k = \emptyset$, where for each $0 < i < k$, each vertex in A_{i-1} is chosen independently to be in A_i with probability $n^{-1/k}$. A point $z \in A_i$ is called an i-pivot of v if $d_G(v, z) = d_G(v, A_i)$. The cluster of a vertex $u \in A_i \setminus A_{i+1}$ is defined as

$$C(u) = \{v \in V \ : \ d_G(u, v) < d_G(v, A_{i+1})\} \ . \tag{21}$$

It was shown in [TZ01a] that

Claim 9. *With high probability, each vertex is contained in at most $4n^{1/k} \log n$ clusters.*

We recall a few definitions from [EN16b]. For each $v \in V$ and $0 \leq i \leq k-1$, a point $\hat{z} \in A_i$ is called an *approximate i-pivot* of v if

$$d_G(v, \hat{z}) \leq (1+\epsilon) d_G(v, A_i) \ . \tag{22}$$

Define

$$C_\epsilon(u) = \{v \in V \ : \ d_G(u, v) < \frac{d_G(v, A_{i+1})}{1+\epsilon}\}. \tag{23}$$

The *approximate cluster* $\tilde{C}(u)$ will be any set that satisfies the following:

$$C_{6\epsilon}(u) \subseteq \tilde{C}(u) \subseteq C(u) \ . \tag{24}$$

It was shown in [EN16b] that once we obtain approximate clusters as trees of G, with $\epsilon \leq 1/(48k^4)$, and provide a routing scheme for these trees, it implies a routing scheme for G with stretch $4k-5+o(1)$. In fact, it suffices that the routing scheme for each tree always routes through the root of the tree, not necessarily via the shortest path in the tree.

Let $h(u,v)$ denote the number of vertices on the shortest path from u to v in G. The following were also shown in [EN16b] to hold with high probability.[8]

Claim 10. *For any $u, v \in V$ with $h(u,v) \geq B$, there exists a vertex of $A_{k/2}$ on the shortest path between them.*

Claim 11. *For any $0 \leq i < k-1$, $v \in A_i \setminus A_{i+1}$ and $u \in C(v)$, it holds that $h(u,v) \leq 4n^{(i+1)/k} \ln n$.*

In particular, for $i < k/2$ we can find the "exact" cluster $C(v)$ for each $v \in A_i \setminus A_{i+1}$, by a simple limited Bellman-Ford exploration from all such vertices v to hop-depth $4n^{(i+1)/k} \ln n \leq \tilde{O}(\sqrt{n})$. By Claim 9, the congestion induced at each $u \in V$ by the merit of being a part of many clusters is only $4n^{1/k} \log n$. So the total number of rounds required is $\tilde{O}(n^{1/2+1/k})$, and each vertex needs to store at most $4n^{1/k} \log n$ words (the clusters containing it). Finally, note that these clusters indeed correspond to trees, since every vertex $u \in C(v)$ can store as a parent the vertex who last updated the distance estimate that u has for v.

From now on we consider the high levels, where $i \geq k/2$. Define $G' = (V', E')$ as a virtual graph where $V' = A_{k/2}$, and E' corresponds to B-bounded distances in G. Observe that Claim 10 implies that $d_{G'}(v,v') = d_G(v,v')$ for any $v, v' \in V'$ (because any shortest path in G has a vertex of V' within any B hops on that path). First, we compute a (β,ϵ)-hopset H for the virtual graph G' as in Theorem 7, with parameters $\log k$, ϵ and $\rho = 1/k$. If one desires the second assertion of the Theorem 10, pick $\rho = \sqrt{\log\log n / \log n}$. Note that the graph G' is implicit, and every node has internal memory $\tilde{O}(m^\rho)$. Since $|A_{k/2}| \leq O(\sqrt{n})$ whp, the number of rounds required to compute H is at most $(n^{1/2+1/k}+D)\cdot(\log n)^{O(1/\rho)}$ (recall $\rho \geq 1/k$ and $\epsilon \geq \Omega(1/\log^4 n)$).

Approximate Pivots To compute the approximate pivots, conduct a Bellman-Ford exploration to depth β in $G'' = G' \cup H$, as in Lemma 8, rooted in A_{i+1}, to compute for each $v \in V'$ a value $\hat{d}(v, A_{i+1})$. We perform another B-bounded exploration in G, where initially every vertex $v \in V'$ sends its current estimate, and in every step every vertex forwards the smallest value it has heard so far. We claim that every $u \in V$ will learn of an approximate $(i+1)$-pivot $\hat{z} \in A_{i+1}$. To see this, let z be the $(i+1)$-pivot of u. If $h(u,z) \leq B$, then u will hear z's message in the last B-bounded exploration. Otherwise, by Claim 10, there exists a vertex $v' \in V'$ on the shortest path from u to z within B hops from u, and since H is a (β,ϵ)-hopset, we have that the first β rounds of Bellman-Ford exploration from A_{i+1} caused v' to update $\hat{d}(v', A_{i+1}) \leq (1+\epsilon) d_G(v', A_{i+1})$. In the final exploration to range B, the vertex v' will communicate this value on the path towards u. Thus, u will have a value at most

$$\hat{d}(u, A_{i+1}) \leq d_G^{(B)}(u,v') + \hat{d}(v', A_{i+1}) \leq d_G(u,v') + (1+\epsilon) d_G(v', A_{i+1}) \leq (1+\epsilon) d_G(u, A_{i+1}) \ , \tag{25}$$

where the last inequality used that $d_G(u,v') + d_G(v', A_{i+1}) = d_G(u, A_{i+1})$. This follows since v' lies on the shortest path from u to the nearest vertex of A_{i+1}. We conclude that no matter which $\hat{z}$ is the approximate pivot of u, the distance estimate that u has for it cannot be larger than $(1+\epsilon) d_G(u, A_{i+1})$. Computing the approximate pivots requires $\tilde{O}(m^{1+\rho} + D) \cdot \beta = (n^{1/2+1/k} + D) \cdot (\log n)^{O(1/\rho)}$ rounds.

[8] For the sake of simplicity we will assume k is even. For odd k, we can improve the running time by a factor of $n^{1/(2k)}$.

Approximate Clusters Fix some $i \geq k/2$, and for each $v \in A_i \setminus A_{i+1}$ we conduct a *limited* Bellman-Ford exploration in $G'' = G' \cup H$ for β rounds rooted at v, as in Lemma 8. By "limited", we mean that any vertex $u \in V'$ receiving a message originated at v, will forward it to its neighbors iff the current distance estimate is strictly less than $\hat{d}(u, A_{i+1})/(1+\epsilon)^2$. We will refer to this condition, the *inclusion condition* of the exploration of v. We need to avoid congestion at intermediate vertices during the B-bounded exploration in G described in Lemma 8, so these vertices will also need to implement some sort of limitation. Concretely, vertices $u \in V \setminus V'$ will forward v's message iff their current estimate is strictly less than $\hat{d}(u, A_{i+1})/(1+\epsilon)$. The exploration over edges of H is done as before, where Claim 9 guarantees every vertex participates in $4n^{1/k} \log n$ clusters (we will soon show that the approximate clusters are indeed contained in the clusters), so this bounds the number of rounds required by $\tilde{O}(n^{1/2+1/k} + D) \cdot \beta$. Also the memory per vertex required from this computation is bounded by $\tilde{O}(n^{1/k})$ (the number of cluster containing the vertex).

This exploration constructs a virtual tree rooted at v. For every edge $(x, y) \in E'$ on this tree, we add to the cluster all the vertices in G on the B-bounded path from x to y. This can be done via an acknowledgement message from y back to x on this path, and every vertex updates its parent accordingly. For every hopset edge (x, y) of the tree (which was broadcast to the entire graph during the exploration), every vertex $u \in P_{x,y}$, where $P_{x,y}$ is the path in G implementing the edge (x, y), joins the tree (u knows about being a part of this edge by the path-reporting property of our hopset), and sets its distance estimate as $b_v(x) + d_P(x, u)$ if this value is smaller than its current estimate. If this is the case, the vertex u also sets its parent as the neighbor on $P_{x,y}$ which is closer to x.

Finally, we perform another limited Bellman-Ford exploration to depth B in G, where every vertex in the tree of v sends its current distance estimate, and every vertex $u \in V$ will forward the smallest estimate it heard so far, but iff it is strictly less than $\hat{d}(u, A_{i+1})/(1+\epsilon)$. In that case it will also join the approximate cluster of v, and will update its parent as its neighbor in G whose message caused u to update its distance estimate to v for the last time.

Observe that the same vertex may join a tree more than once, due to several edges in $E' \cup H$ whose paths contain it. In such a case the vertex will have as a parent the vertex which minimize the estimated distance to the root. Since every vertex has a single parent, we will have that the approximate cluster of v, $\tilde{C}(v)$, is indeed a tree. It remains to prove (24). Let $b_v(u)$ be the distance estimate that u has to v in the exploration rooted at v.

Claim 12. *For any $v \in V'$, $\tilde{C}(v) \subseteq C(v)$.*

Proof. Consider any $u \in \tilde{C}(v)$. If it is the case that $u \in V$ joined the approximate cluster by the exploration rooted at v, either by being in V' or on a B-bounded path in G that implements an edge of E', then it must satisfy $b_v(u) < \hat{d}(u, A_{i+1})/(1+\epsilon)$. Now,

$$d_G(u, v) \leq b_v(u) < \hat{d}(u, A_{i+1})/(1+\epsilon) \overset{(25)}{\leq} d_G(u, A_{i+1}) \ ,$$

so indeed $u \in C(v)$. The other case is that $u \ in P_{x,y}$ for a path $P_{x,y}$ implementing a hopset edge (x, y) that was added to the virtual tree. Since y joins the approximate cluster, it must satisfy $b_v(y) < \hat{d}(y, A_{i+1})/(1+\epsilon)^2$. Recall that the weight of the hopset edge $w_H(x, y)$ is the weight of the path $P = P_{x,y}$ from x to y in G

that u lies on. Hence $d_P(x,u)+d_P(u,y)=w_H(x,y)$. It follows that

$$\begin{aligned} d_G(u,A_{i+1}) &\overset{(25)}{\geq} \frac{\hat{d}(u,A_{i+1})}{1+\epsilon} \geq \frac{d_G(u,A_{i+1})}{1+\epsilon} \geq \frac{d_G(y,A_{i+1})-d_G(u,y)}{1+\epsilon} \\ &\overset{(25)}{\geq} \frac{\hat{d}(y,A_{i+1})}{(1+\epsilon)^2} - \frac{d_P(u,y)}{1+\epsilon} > b_v(y) - d_P(u,y) \\ &= b_v(x)+w_H(x,y)-d_P(u,y) = b_v(x)+d_P(x,u) \\ &\geq b_v(u) \geq d_G(u,v)\ , \end{aligned}$$

where in the penultimate inequality we used the fact that the vertex u knows $d_P(x,u)$, and thus it could have updated its distance estimate to v as $b_v(x)+d_P(x,u)$ (note that it may have used a smaller estimate). Thus $u\in C(v)$ in this case, as required. □

The next claim proves the second inequality of (24).

Claim 13. *For any $v\in V'$, $C_{6\epsilon}(v)\subseteq\tilde{C}(v)$.*

Proof. Let $u\in C_{6\epsilon}(v)$. We would like to show that $u\in\tilde{C}(v)$. Consider the shortest path P from u to v in G. Then by Claim 10, there is a vertex $u'\in V'$ on P that is within B hops from u. Notice that

$$d_G(v,u')=d_G(v,u)-d_G(u,u')\leq\frac{d_G(u,A_{i+1})-d_G(u,u')}{1+6\epsilon}\leq\frac{d_G(u',A_{i+1})}{1+6\epsilon}\ . \tag{26}$$

Hence $u'\in C_{6\epsilon}(v)$ too.

We will show that the limited exploration originated at v will reach u', and in the final depth B exploration it will reach u and include it in $\tilde{C}(v)$.

Since H is a (β,ϵ)-hopset, there is a path P' in G'' from v to u' that contains at most β edges that satisfies

$$d_{P'}(v,u')\leq(1+\epsilon)d_{G'}(v,u')=(1+\epsilon)d_G(v,u')\ . \tag{27}$$

Let $z\in P'$ be any vertex on P' that lies t hops from v, $0\leq t\leq\beta$. Then after t steps of Bellman-Ford exploration from v we have that

$$\begin{aligned} b_z(v) &= d_{P'}(v,z)=d_{P'}(v,u')-d_{P'}(z,u') \\ &\overset{(27)}{\leq} (1+\epsilon)d_G(v,u')-d_G(z,u') \overset{(26)}{\leq} \frac{(1+\epsilon)d_G(u',A_{i+1})}{1+6\epsilon}-d_G(z,u') \\ &\leq \frac{d_G(u',A_{i+1})-d_G(z,u')}{1+4\epsilon} < \frac{d_G(z,A_{i+1})}{(1+\epsilon)^2} \leq \frac{\hat{d}(z,A_{i+1})}{(1+\epsilon)^2}\ . \end{aligned}$$

(We used that $\epsilon<1/5$.) We conclude that z satisfies the inclusion condition for the exploration rooted at v, and forwards the message of v onwards. In particular, by (27), $b_v(u')\leq d_{P'}(v,u')\leq(1+\epsilon)d_G(v,u')$. In the final phase we make a Bellman-Ford exploration for B rounds in G from each vertex that received the message of v. Thus, u' will start such an exploration with distance estimate $b_v(u')$. Consider the subpath $Q\subseteq P$ from u' to u. We have to show that every vertex on this path forwards the message of v, that is, that it satisfies the inclusion condition of the exploration of v. Let $y\in Q$ be such a vertex. Since this is a shortest

path in G, we have

$$\begin{aligned} b_v(y) &\leq b_v(u') + d_Q(u', y) \leq (1+\epsilon) d_G(v, u') + d_G(u', y) \\ &\leq (1+\epsilon) d_G(v, y) = (1+\epsilon)(d_G(v, u) - d_G(y, u)) \\ &\overset{(23)}{\leq} \frac{(1+\epsilon) d_G(u, A_{i+1})}{1+6\epsilon} - d_G(y, u) \leq \frac{d_G(u, A_{i+1}) - d_G(y, u)}{1+4\epsilon} \\ &\leq \frac{d_G(y, A_{i+1})}{1+4\epsilon} < \frac{\hat{d}(y, A_{i+1})}{1+\epsilon}\,, \end{aligned}$$

as required.

□

5.1 Distributed Tree Routing with Small Memory

In this section we present our compact routing scheme for trees that can be computed in a distributed manner using small internal memory. In previous constructions of distributed routing schemes for trees [EN16b, LPP16], the internal memory was as high as $\sqrt{n}$, and it was also somewhat inefficient: the label size is $O(\log^2 n)$ and the routing tables are of size $O(\log n)$. Compare this to the classical [TZ01b] tree routing, which has label size $O(\log n)$ and routing tables of size $O(1)$.

We follow the basic framework of previous works, by selecting a set $U \subseteq V$, such that each vertex is sampled to U independently with probability q (q is a parameter, which we shall optimize later). Fix a tree T on vertices $V(T) \subseteq V$ with root z. The vertices $U(T) = (U \cap V(T)) \cup \{z\}$ partition the tree into subtrees, by removing the edges from each vertex in $U(T)$ to its parent. Each of the $|U(T)|$ subtrees is rooted in a vertex of $U(T)$. Denote by T_w the subtree rooted at w. We also consider T', the virtual tree on the vertices of $U(T)$, which is rooted at z, and contains an edge (x, y) if the parent of y lies in T_x. It is not hard to see (e.g., [EN16b]) that whp the depth of each T_w is $\tilde{O}(1/q)$, and that $|U| \leq O(qn)$.

In both [EN16b, LPP16], routing schemes were created for each T_w, and also a routing scheme for the virtual tree T'. This computation required large internal memory, since z had to locally compute the scheme for T'. The inefficiency in the size was due to the fact that when routing in T', traveling over a virtual edge (x, y), one has to route in T_x from x to the parent of y. This seems to require storing additional routing information for this subtree, increasing both label and table size by a logarithmic factor. We overcome this issue by storing routing information only with respect to the actual tree, while applying pointer jumping techniques to quickly compute the full labels. However, we do not know how to construct exact tree routing with small memory. Fortunately, to implement our routing scheme for general graphs, it suffices to provide a *root-tree routing* scheme, where the routing is always done via the root of the tree T, and not necessarily via the shortest path. (We stress that using larger memory, we can compute exact tree routing tables and labels within $\tilde{O}(\sqrt{n} + D)$ rounds, with label size $O(\log n)$ and routing tables of size $O(1)$, substantially improving previous results.)

Before describing our approach, let us briefly recall the Thorup-Zwick construction of tree routing. The idea is to assign to every (non-leaf) vertex $x \in T$ its *heavy child*, which is the child whose subtree has maximal size. Note that the subtree of any non-heavy child of x contains at most half of the vertices of the subtree T_x of T rooted at x. For this reason, any path from the root z to some $y \in T$ contains at most $\log n$ non-heavy edges. For an exact routing scheme they also conducts a DFS search in T that assigns to each y the DFS entry and exit times for its subtree. The label of y is these entry and exit times, and also the names of the non-heavy edges on the z to y path. The routing table y consists of the DFS times, the name of the heavy child, and the name of the parent of y in the tree. The routing towards a target v in the tree is done

as follows. At any intermediate vertex $y \in T$, if v is not in the subtree rooted at y (this can be checked via the DFS times), then y forwards to its parent. If v is in the subtree, y inspects v's label to see if an edge (y, x) appears there. If this is the case, it forwards to x, otherwise to its heavy child. Note that if one desires root-tree routing then there is no need to implement a DFS – initially route to the parent until the root is reached, and then follow the path using heavy edges unless the label indicates otherwise.

Now we show how to implement our scheme in a distributed manner, and with $O(\log n)$ internal memory. First, every $w \in U(T)$ sends a message about itself to the vertices of T_w, informing them they are in T_w. Note that this message will arrive to all vertices in $T(U)$ who are children of w in the virtual tree T', so they will know their parent. Next, for each $w \in U(T)$, every vertex in T_w sends to its parent the size of the subtree rooted at it, beginning with the leaves. Every vertex that received messages from all its children, sums up the values and sends to its own parent. This can be done in parallel for all trees T_w for $w \in U(T)$, and will take $\tilde{O}(1/q)$ (the bound on the height of each T_w) rounds.

For a vertex v in a tree T, rooted at a vertex z, and a positive integer h, we say that a vertex u is an *h-ancestor* of v, if u lies on the unique $v - z$ path in T at distance h from v.

We would like that every $y \in T$ will know the entire size of the subtree of T rooted at y. Initially, we compute this value only for the virtual vertices of $U(T)$. For a vertex $x \in U(T)$, its subtree size is exactly the sum of sizes of subtrees T_w for w that are in the subtree of T' rooted at x. Note that computing these values from the leaves of T' up will not be efficient, since every message on a virtual edge may require $O(D)$ rounds, and the depth of T' may be as large as qn (which will be approximately $\sqrt{n}$). Thus, this results in $O(D\sqrt{n})$ rounds. To alleviate this issue, we use the following "pointer jumping" technique. Initially, set for $x \in U(T)$ the current size $s_x = |T_x|$, and its first ancestor $a_1(x)$ as its parent in T' (and for the root z, set $a_1(z) = \perp$). For $i = 0, 1, \ldots, \log n$ rounds, every vertex $x \in U(T)$ will broadcast in the ith round (using the BFS tree of G), the current size s_x and the name of its 2^i-ancestor $a_i(x)$ in T'. Then whenever x hears a message that some $w \in U(T)$ broadcasts with $x = a_i(w)$, then x adds s_w to its current size s_x. In addition, the vertex x hears the message of $a_i(x)$, and it updates $a_{i+1}(x)$ as $a_i(a_i(x))$. (It could be the case that $a_i(a_i(x)) = \perp$. In this case, indeed, $a_{i+1}(x) = \perp$.) We claim that this process correctly computes for any $x \in U(T)$ the size of the subtree of T rooted at x. It can be shown by induction on i, that before the ith round, s_x is the size of the subtree rooted at x that contains at most 2^i vertices of $U(T)$ on any root-leaf path. There are $O(|U(T)|) \leq \tilde{O}(qn)$ messages sent on each round for $\log n$ rounds. Hence, it will take $\tilde{O}(qn + D)$ rounds to implement this step.

In order to compute s_y, the size of the subtree of T rooted at y, for all $y \in T$, every $x \in U(T)$ informs its parent in T with the value s_x. Then once again, for every $w \in U(T)$ in parallel, the leaves of T_w start to send to their parent their current size. This time, some of these leaves and internal vertices could be parents of vertices in $U(T)$, so these sizes are the actual subtree size in T. In $\tilde{O}(1/q)$ rounds, every vertex $y \in T$ will know s_y. After sending these values to the parents, every vertex can infer who is its heavy child.

The label $L(y)$ needed for root-tree routing is just the collection of edges $\{(u, v)\}$ that are on the $z - y$ path in T, such that v is not the heavy child of u. Clearly, there can be at most $\log n$ such edges on this path, because the size of the subtree decreases by a factor of 2 for every non-heavy edge. If $y \in T_x$, we start by computing a partial label that contains non-heavy edges on the path from x to y. This can be done by initializing $L(x) = \emptyset$, and starting at x, any vertex $u \in T_x$ which received a label $L(u)$, sends $L(u)$ to its heavy child, and $L(u) \cup \{(u, v)\}$ for any non-heavy child v. These labels are also sent to the children of x in T' (recall that these are the vertices $T(U)$ whose T-parents belong to T_x). Once this computation is completed, every vertex $w \in T(U)$ knows the non-heavy edges on the path from x, its parent in T', to w. We again apply pointer jumping to compute the full labels. For $i = 0, 1, \ldots, \log n$, every vertex of $U(T)$ will broadcast in the ith round its current label. In each round, when x hears the message from its 2^j-ancestor

$a_j(x)$ (recall that x computed previously its 2^j-ancestors, for all $j = 0, 1, \ldots, \log n$, and it stored them in its internal memory), it will update $L(x) \leftarrow L(a_j(x)) \circ L(x)$. Once again, it can be proved by induction on i that before the ith round, every $x \in U(T)$ knows all the non-heavy edges on the path in T from $a_i(x)$ to x (or from the root z to x if $a_i(x) = \perp$). Since every label has size $O(\log n)$, this will require $\tilde{O}(qn + D)$ rounds. Finally, in another $\tilde{O}(1/q)$ rounds, each $x \in U(T)$ sends its updated label $L(x)$ to every vertex $y \in T_x$, and they update their label by appending $L(x)$.

If one desires a routing scheme for a single tree, just take $q = 1/\sqrt{n}$, so the running time will be $\tilde{O}(\sqrt{n} + D)$. If we desire to compute a routing scheme in parallel for multiple trees, but have the guarantee that every $v \in V$ belongs to at most s trees, then we can use the argument as in [EN16a] to obtain running time $\tilde{O}(\sqrt{s \cdot n} + D)$ (rather than the naive $\tilde{O}(s \cdot \sqrt{n} + D)$). We conclude by formally summarizing our result.

Theorem 11. *For any tree T on n vertices, lying in a network with hop-diameter D, there exists a distributed algorithm in the CONGEST model running in $\tilde{O}(\sqrt{n}+D)$ rounds, that computes a root-tree routing scheme with label size $O(\log n)$ and routing tables of size $O(1)$, such that every vertex uses only $O(\log n)$ words of memory throughout the computation.*

Moreover, if there are no restriction on the memory used throughout the computation, then exact tree routing tables of size $O(1)$ and labels of size $O(\log n)$ can be computed in $\tilde{O}(\sqrt{n} + D)$ time.

In addition, given a network with n vertices and a set of trees so that each vertex is contained in at most s trees, one can compute a root-tree routing scheme as above for all trees in parallel, within $\tilde{O}(\sqrt{s \cdot n} + D)$ rounds, while using memory $O(s \cdot \log n)$ at each vertex.

Acknowledgements

We wish to thank Christoph Lenzen for raising to us the problem of distributed routing with small individual memory requirements, and for permitting us to use a quotation from [Len16].

References

[ABCP93] Baruch Awerbuch, Bonnie Berger, Lenore Cowen, and David Peleg. Near-linear cost sequential and distributed constructions of sparse neighborhood covers. In *34th Annual Symposium on Foundations of Computer Science, Palo Alto, California, USA, 3-5 November 1993*, pages 638–647, 1993.

[ABNLP90] Baruch Awerbuch, Amotz Bar-Noy, Nathan Linial, and David Peleg. Improved routing strategies with succinct tables. *J. Algorithms*, 11(3):307–341, September 1990.

[ABP17] Amir Abboud, Greg Bodwin, and Seth Pettie. A hierarchy of lower bounds for sublinear additive spanners. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 568–576, 2017.

[AGM04] Ittai Abraham, Cyril Gavoille, and Dahlia Malkhi. Routing with improved communication-space trade-off. In *Distributed Computing, 18th International Conference, DISC 2004, Amsterdam, The Netherlands, October 4-7, 2004, Proceedings*, pages 305–319, 2004.

[AGM+08] Ittai Abraham, Cyril Gavoille, Dahlia Malkhi, Noam Nisan, and Mikkel Thorup. Compact name-independent routing with minimum stretch. *ACM Trans. Algorithms*, 4(3):37:1–37:12, July 2008.

[AP92] B. Awerbuch and D. Peleg. Routing with polynomial communication-space tradeoff. *SIAM J. Discrete Mathematics*, 5:151–162, 1992.

[Ber09] Aaron Bernstein. Fully dynamic (2 + epsilon) approximate all-pairs shortest paths with fast query and close to linear update time. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*, pages 693–702, 2009.

[Che13] Shiri Chechik. Compact routing schemes with improved stretch. In *ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013*, pages 33–41, 2013.

[Coh93] Edith Cohen. Fast algorithms for constructing t-spanners and paths with stretch t. In *34th Annual Symposium on Foundations of Computer Science, Palo Alto, California, USA, 3-5 November 1993*, pages 648–658, 1993.

[Coh97] Edith Cohen. Using selective path-doubling for parallel shortest-path computations. *J. Algorithms*, 22(1):30–56, 1997.

[Coh00] Edith Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. *J. ACM*, 47(1):132–166, 2000.

[Cow01] Lenore Cowen. Compact routing with minimum stretch. *J. Algorithms*, 38(1):170–183, 2001.

[EGP03] Tamar Eilam, Cyril Gavoille, and David Peleg. Compact routing schemes with low stretch factor. *J. Algorithms*, 46(2):97–114, 2003.

[Elk17] Michael Elkin. Distributed exact shortest paths in sublinear time. 2017. To appear in STOC'17.

[EN16a] Michael Elkin and Ofer Neiman. Hopsets with constant hopbound, and applications to approximate shortest paths. In *57th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2016, New Brunswick, NJ, USA, October 2016 (to appear)*, 2016.

[EN16b] Michael Elkin and Ofer Neiman. On efficient distributed construction of near optimal routing schemes: Extended abstract. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, PODC '16, pages 235–244, New York, NY, USA, 2016. ACM.

[EP04] Michael Elkin and David Peleg. (1+epsilon, beta)-spanner constructions for general graphs. *SIAM J. Comput.*, 33(3):608–631, 2004.

[GGHI13] Cyril Gavoille, Christian Glacet, Nicolas Hanusse, and David Ilcinkas. On the communication complexity of distributed name-independent routing schemes. In *Distributed Computing - 27th International Symposium, DISC 2013, Jerusalem, Israel, October 14-18, 2013. Proceedings*, pages 418–432, 2013.

[GP03] Cyril Gavoille and David Peleg. Compact and localized distributed data structures. *Distributed Computing*, 16(2-3):111–120, 2003.

[HKN14] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Decremental single-source shortest paths on undirected graphs in near-linear total update time. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 146–155, 2014.

[HKN16] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A deterministic almost-tight distributed algorithm for approximating single-source shortest paths. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 489–498, New York, NY, USA, 2016. ACM.

[KS97] Philip N. Klein and Sairam Subramanian. A randomized parallel algorithm for single-source shortest paths. *J. Algorithms*, 25(2):205–220, 1997.

[Len16] Christoph Lenzen. Personal communication, 2016.

[LP13] Christoph Lenzen and Boaz Patt-Shamir. Fast routing table construction using small messages. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 381–390, 2013.

[LP15] Christoph Lenzen and Boaz Patt-Shamir. Fast partial distance estimation and applications. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21 - 23, 2015*, pages 153–162, 2015.

[LPP16] Christoph Lenzen, Boaz Patt-Shamir, and David Peleg. Distributed distance computation and routing with small messages, 2016. manuscript.

[MPVX15] Gary L. Miller, Richard Peng, Adrian Vladu, and Shen Chen Xu. Improved parallel algorithms for spanners and hopsets. In *Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '15, pages 192–201, New York, NY, USA, 2015. ACM.

[Nan14] Danupon Nanongkai. Distributed approximation algorithms for weighted shortest paths. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 565–573, 2014.

[Pel00] David Peleg. *Distributed Computing: A Locality-sensitive Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.

[PU89] David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *J. ACM*, 36(3):510–530, 1989.

[SHK+12] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. Distributed verification and hardness of distributed approximation. *SIAM J. Comput.*, 41(5):1235–1265, 2012.

[SS99] Hanmao Shi and Thomas H. Spencer. Time-work tradeoffs of the single-source shortest paths problem. *J. Algorithms*, 30(1):19–32, 1999.

[TZ01a] M. Thorup and U. Zwick. Approximate distance oracles. In *Proc. of the 33rd ACM Symp. on Theory of Computing*, pages 183–192, 2001.

[TZ01b] Mikkel Thorup and Uri Zwick. Compact routing schemes. In *Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures*, SPAA '01, pages 1–10, New York, NY, USA, 2001. ACM.

[TZ06] M. Thorup and U. Zwick. Spanners and emulators with sublinear distance errors. In *Proc. of Symp. on Discr. Algorithms*, pages 802–809, 2006.

[UY91] Jeffrey D. Ullman and Mihalis Yannakakis. High-probability parallel transitive-closure algorithms. *SIAM J. Comput.*, 20(1):100–125, 1991.